

AFRL-IF-RS-TR-2004-139
Final Technical Report
June 2004



SECURE ARCHITECTURE FOR EXTENSIBLE MOBILE INTERNET TRANSPORT SYSTEM

State University of New York Institute of Technology at Utica-Rome

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-139 has been reviewed and is approved for publication

APPROVED: /s/

ANDREW J. KARAM
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE JUNE 2004	3. REPORT TYPE AND DATES COVERED Final Apr 01 – Apr 02	
4. TITLE AND SUBTITLE SECURE ARCHITECTURE FOR EXTENSIBLE MOBILE INTERNET TRANSPORT SYSTEM			5. FUNDING NUMBERS C - F30602-01-1-0518 PE - 62702F PR - OIAG TA - 32 WU - P3	
6. AUTHOR(S) Digen Das, Patrick Fitzgibbons, and Larry Hash				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) State University of New York Institute of Technology at Utica-Rome Rome New York 13440			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFGB 525 Brooks Road Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-139	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Andrew J. Karam/IFGB/(315) 330-7290/ Andrew.Karam@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This document is the final report for the State University of New York Institute of Technology Secure Architecture For Extensible Mobile Internet Transport System project. In this paper, we will summarize our accomplishments, discuss high and low points of the project, and suggest future work that might be done to further mobile/wireless security research. We will also present a final overview discussion of Mobile Security Policy. Although our initial AFRL funding period is over we intend to pursue more secure wireless research, and will use this site to post any future results in software or technical reports. We would like to thank AFRL for giving us the opportunity to get started in this research arena.				
14. SUBJECT TERMS Secure Architecture, Extensible Mobile Internet, Wireless Security, Mobile Wireless Security, Mobile Security Policy			15. NUMBER OF PAGES 85	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1 Introduction-----	1
2 Accomplishments -----	1
2.1 Creation of a secure enclave model for wireless mobility that includes both inter and intra-domain Mobile-IP-----	1
2.2 Integration of Mobile-IP and IPSEC in terms of routing and security -----	1
2.3 Simplified Link-Layer Only Ad Hoc Routing-----	2
2.4 The Establishment of Wireless Campus Infrastructures-----	3
2.5 Availability of Wireless drivers for LAPTOP and DESKTOP platforms-----	3
3 High Points-----	3
3.1 Combined Mobile-IP and IPSEC -----	3
3.2 Interest in wireless systems by IT administrators at SUNY Institute-----	4
4 Low Points -----	4
4.1 Death By Integration -----	4
4.2 Wireless a Moving Target -----	4
5 Mobile Security Policy Overview-----	5
5.1 Secure Enclave approach -----	5
5.2 Us versus Them -----	6
5.3 Foreign Agent Considerations-----	7
5.4 Home Agent Considerations-----	7
5.5 Mobile Node Considerations-----	8
6 Suggested Further Work-----	9
6.1 Mobile Nodes Abroad -----	9
6.2 Smarter Foreign Agents -----	10
6.3 The hard work – integration-----	10
6.4 Keys as a basis for networking-----	10
6.5 Wireless Loading -----	11
7 Acknowledgements -----	11
Appendix A Secure Architecture for Extensible Mobile Internet Transport Systems Seventh Quarterly Report-----	12
Appendix B A Secure Architecture for Extensible Mobile Internet Transport Systems Tenth Quarterly Report -----	21
Appendix C A Secure Architecture for Extensible Mobile Internet Transport Systems Ninth Quarterly Report-----	29
Appendix D A Secure Architecture for Extensible Mobile Internet Transport Systems Eighth Quarterly Report -----	39
Appendix E A Secure Architecture for Extensible Mobile Internet Transport Systems Seventh Quarterly Report-----	55

Appendix F A Secure Architecture for Extensible Mobile Internet Transport Systems Sixth Quarterly Report -----	65
Appendix G A Secure Architecture for Extensible Mobile Internet Transport Systems First Quarterly Report -----	75
BIBLIOGRAPHY -----	80

Section 1: Introduction

This document is the final report for the State University of New York Institute of Technology Secure Architecture For Extensible Mobile Internet Transport System project. In this paper, we will summarize our accomplishments, discuss high and low points of the project, and suggest further work that might be done to further mobile/wireless security research. We will also present a final overview discussion of Mobile Security Policy. Although our initial AFRL funding period is over, we intend to pursue more secure wireless research, and will use this site to post any future results in software or technical reports. We would like to thank AFRL for giving us the opportunity to get started in this research arena.

Section 2: Accomplishments

In this section, we will briefly present our accomplishments (aspects which were not accomplished will be relegated to the low points section below). We will minimize details and provide only bare bones summary text. Please note that Appendix A below serves to tie these accomplishments to our quarterly reports. We hope that the appendix may serve as a bibliographic guide to content in previous reports. Our accomplishments are discussed in the following subsections.

2.1 Creation of a secure enclave model for wireless mobility that includes both inter and intra-domain Mobile-IP.

We constructed an integrated Mobile-IP/IPSEC system in which Mobile Nodes abroad can use 2-way tunnels to securely tunnel all their packets to and from their Home Agent (assumed to be at home in the secure enclave, of course). This solution only addresses one possible facet of a many-sided security problem. We also invented two forms of ad hoc routing (including multi-hop) and tied them to end-to-end (but network-layer) IPSEC-based routing. Thus hosts that a priori belong to the same security enclave may choose to securely talk to their security peers. Further our Home Agent and Foreign Agents use one-way tunnels authenticated with AH. This allows all agents to reject any arriving (tunnel) packets that do not have an IPSEC binding using agent IP addresses.

Our 2-way tunnels deal with the problem of what to do about "our" mobile nodes, but neglect dealing with non-local mobile nodes. Issues here are complicated and we cannot claim to have the last word. However we have addressed many security issues in this area as described in Internet draft [2] and refer interested readers to that document. In brief, we suggest that foreign nodes simply be logically treated as being "outside" the enclave and their packets need only be tunneled across the enclave to a firewall access point.

2.2 Integration of Mobile-IP and IPSEC in terms of routing and security.

A key focus of our work was to tie Mobile-IP and IPSEC directly together. At this point in time, Cisco (for example) has made both Mobile-IP and IPSEC available in their routing

devices in IOS version 12.0 [5]. However there is no evidence that the two have been integrated. We believe that the combination of IPSEC and Mobile-IP is far superior to virtual link-layer tunneling schemes such as L2TP [6] or Microsoft PPTP [7] simply because IPSEC has wide-spread architectural utility, generality, and is liable to be more supported than more proprietary schemes. For example, IPSEC can cover both the link layer (by being at the network layer) or the transport layer or run router to end system, router to router, etc. There is also no point in separate security mechanisms for the latest virtual link tunneling scheme when IPSEC can be used in all places. Mobile-IP needs IPSEC by definition as all packets between a remote Mobile Node (or Mobile Node on a wireless link) to/from home should be made secure.

In our system, we tied IPSEC to routes. For example, when a Mobile Node installs a default route, it is aiming that route at an agent. A route binding that included an indirection mechanism (tunnel IPSEC to the Home Agent) was part of the picture. When we did non Mobile-IP work (say using our multi-hop ad hoc routing protocol) to tie two Mobile Nodes together, we also naturally tied IPSEC to the host routes installed in each Mobile Node. All other IPSEC implementations we have seen so far tie IPSEC to some sort of additional packet-filter like access list mechanism. Our approach seems more powerful and in some ways simpler, but to be fair we cannot make a compelling case for its superiority over access list mechanisms (other than one less lookup in the IP layer, but even that is not terribly important given the relatively blinding speed of processors these days). Our IPSEC route binding mechanism can however easily be used to setup manual Virtual Private Networks between two routes simply by installing symmetric keys in a key file, and using the route(8) administrative command to install a route (to a network, subnet, or host). Our mechanism also applies to ARP/link-layer bindings. Our architecture seems to have general utility.

2.3 Simplified Link-Layer Only Ad Hoc Routing

In this protocol, we do not use ARP on a link. We instead use a protocol (like the ISO ES-IS) in which all nodes, agents, and Mobile Nodes send authenticated beacons. This "non" ARP mechanism is intended to serve a number of purposes. First, it tries to mitigate possible ARP spoofing by insisting that the (IP address, MAC address) binding be authenticated. Note that we have implemented this mechanism with both symmetric and asymmetric key systems (in the former case, we have a network-wide key; in the latter, a per host signature). Secondly, the mechanism serves to tie networks together by key possession. It is not important if two laptops do or do not share a subnet. All systems beacon. Therefore if you share a key, you can talk. The low-level IP subnet semantic that requires a router for communication between two hosts from different subnets is obviated. The mechanism also serves a gateway function so that systems, which do not possess the secret, cannot penetrate into a secure enclave through a "firewall-like" mobility agent. Lastly the mechanism serves a very important purpose in that we assume that if we can hear your beacon, we can talk to you. Beacons (unlike ARP) is done at a relatively high rate of speed. If a system disappears, we will use other routing mechanisms to try and find it (and not believe an ARP cache entry that is going to hang around for twenty minutes).

The only substantial criticism that can be made of this system is that if everyone beacons the link itself may be less scalable/useable in terms of throughput. Given that current wireless links cannot support many simultaneous hosts anyway, it is hard to understand how this criticism can be valid. ES-IS originally was criticized along these terms, but the critics apparently did not notice that beacon rates were extremely slow (once per minute for fixed ethernet systems). (See our criticism of wireless loading in the Suggested Future Work section below for more discussion on this topic). Slow beacon rates for Mobile Nodes along with a combined unicast ACK "reply" to agent beacons make the mechanism more scalable. Agents can beacon and Mobile Nodes can simply append their MAC address under Mobile-IP authentication when they use Mobile-IP to register. This takes care of Mobile Nodes that only desire to talk to the wired infrastructure and do not want ad hoc service. Non-mobile Mobile Nodes do not need to beacon very often and can probably slow down their beacon rates (as long as agents do not remove them from the routing state in the agent). Highly Mobile Mobile Nodes need to beacon at higher rates in order to talk to each other.

We suspect the real problem here is wireless loading. If the link itself should be able to tolerate 100 FTP transfers at the same time {one should not worry overly much about 100 nodes sending out 100 byte beacons, even if the beacon rate for all nodes is 1 per second. Of course, this would be too much for WAN wireless systems with small overall bandwidth, but in such cases an IEEE 802.11 [8] style link-layer registration protocol only between agent and Mobile Node can make sense; i.e., one could well neglect the ad hoc function when there are too many nodes in a cell, or one simply doesn't care to talk to anything other than the agent (highly likely in many usage scenarios).

2.4 The Establishment of Wireless Campus Infrastructures

We established a wireless network that has a few users (primary research investigators and a few graduate students) within the School of Information Systems and Engineering Technology department of telecommunications[21]. We intend to maintain this network and extend it where possible. A limited three-node network was established and is still in use by the researchers.

2.5 Availability of wireless LAN drivers for LAPTOP and DESKTOP platforms

There are a number of wireless LAN drivers available for various operating systems, including Unix/FreeBSD, Linux, and Windows.

Section 3: Highpoints

3.1 Combined Mobile-IP and IPSEC system

Given the complexity of many of the sub-systems, we are pleased that our system is in use at the SUNY Institute of Technology Information Assurance lab by a small numbers and that the combined IPSEC/Mobile-IP system itself is also in use.

3.2 Interest in wireless systems by IT administrators at SUNY Institute of Technology

The School of Information Systems Engineering Technology faculty and staff and IT administators at SUNYIT have seemingly become very fond of our mobile system (especially network administrators). From the network administrators' point of view, wireless is a natural out of band access system to the network, and represents a secondary path for troubleshooting. An IT staff person may typically be called to service a "broken" computer, that may simply be broken due to network setup problems. It is extremely useful to be able to bring along a working mobile computer that can then be deployed to troubleshoot the fixed wired computer's problem.

Section 4: Low Points

All efforts of this scope have their share of frustration and other low points. We are no exception. This section discusses some of the things that did not go as well as we would have liked.

4.1 Death By Integration

Our system involved kernel drivers, kernel security code, modifications to routing table code, complex application routing daemons, symmetric and asymmetric key-based systems, including an experimental version of DNS. The integration work included at least two kernel ports including a new release in the middle of the project. We ended up with three servers that needed software replacements or upgrades. Certainly, in general, security code is complex, and requires complex skills especially when combined with fundamental operating system level system's programming. The complexity of our system required ever increasing release testing and integration testing. The worst aspect of the problem was probably the number of moving targets including frequent changes in operating system releases. We cannot blame anyone for that, of course, but it seemed that everytime we tried to catch up, before we could work out the bugs, a new version of Mobile-IP would suddenly appear on the scene.

During the last year, when graduate student involvement and participation in the project became more important to our success we did not have the resources to test the more complex aspects and fully integrate them. We did however make a last ditch attempt to finalize kernel modifications in a set of beta release patches. Our hope is that the latest software release may aid in future ports. Our belief is that complex security systems require long-term commitment and extensive test and integration phases.

4.2 Wireless a Moving Target

802.11 as an IEEE project started in the mid 1990's, which was well before our project started. Perhaps not to be overlooked is the fact that the 802.11b standard only became available when our project was getting underway. In the last year alone, IEEE 802.11b wireless hardware has begun to become widely available. Unfortunately not all of our wireless LAN cards and drivers were interoperable when combining equipment including wireless access points from different manufacturers. For example SMC advertised support for third-party drivers on its web pages. When they released the IEEE code, they expunged all mention of third-party drivers and only provided a limited number of PC OS-centric drivers for the IEEE systems, thus making things difficult and only recently has a linux driver (which still has limited features) been made available [9]. We have already purchased a limited amount of IEEE hardware and intend to begin to move our test systems and wireless infrastructure in that direction. In the process we also transitioned to linux as well away from FreeBSD. The good news is that prices per NIC card have fallen approximately to 1/2 or 1/4 of the prices when the project started. (From \$500 per unit when we started to under \$200 for some (non-Lucent) IEEE cards.

Section 5: Mobile Security Policy Overview

In this section, we will briefly review our overall thoughts on mobile security policy. First we will consider the situation from the point of view of a given secure enclave trying to implement mobility, and then we will review the situation from the point of view of the traditional Mobile-IP architectural elements (Mobile Node, Home Agent, Foreign Agent).

5.1 Secure Enclave approach

By "secure enclave" we mean one set of hosts under a single security administration that is protected by some sort of a priori security mechanisms. For example, the secure enclave may be connected up to the Internet but is protected by one or more firewall systems which might either be of the packet filter or bastion host variety. We do not rule out "defense in depth" for interior hosts. We do however assume that some hosts are exposed to the Internet and others are not. We are interested in hosts that are exposed to the Internet. We are also interested in hosts with wireless link layers, which for the sake of argument, we will assume are more susceptible to attacks like promiscuous mode sniffing. In combining Mobile-IP and IPSEC in Mobile Nodes, we first of all made a number of simplifying assumptions. We assumed that a Mobile Node when at home could maintain a two-way IPSEC tunnel connection between it and its Home Agent. We assumed that a Mobile Node when abroad at either a DHCP link or Foreign Agent link could use two-way IPSEC to tunnel home to the Home Agent. The mobile node would dynamically discover these situations and setup tunnels with appropriate security mechanisms. The Home Agent was thus always a bastion host or security gateway to the secure enclave, via any link on it (wireless or wired). Foreign Agents in our first-cut model only serve as link-layer wireless gateways to a secure enclave and/or may not serve external visitors (but must serve internal wireless users, else they are not of interest). IPSEC associations between Mobile Nodes and Foreign Agents did not exist. A Mobile Node abroad might thus be viewed as

an extension of the home secure enclave. The two-way tunnel to and from home would serve as an umbilical code to extend the umbrella of the secure enclave to the Mobile Node.

In more detail, let us look at the relationship between a Mobile Node and Home Agent with the Mobile Node at home. If the Mobile Node is a wireless node, our IPSEC system gives it a functional equivalent for link-layer security. (If it was wired, the user (or the local security authorities) might disable this function). What is important is that all packets bearing ARP or Mobile-IP itself would be subject to network-layer IPSEC security, which might be more or less good depending on the specific IPSEC implementation and security transforms in use. We replaced ARP with a more secure ad hoc mechanism that simply made traditional ARP spoofing more difficult and made two-way exchanges into the secure enclave (via any agent) difficult as well. Mobile-IP had its own authentication mechanism (and yet others have developed a digital signature replacement scheme for its authentication). There is nothing here to prevent the use of additional security at non-network layers including IPSEC at the transport layer, or transport-equivalent security mechanisms like the Finnish ssh (which we use all the time). This mechanism is however very general. The Mobile Node at home is simply equivalent to any current system using its default router. What is curious is that one still finds weak mechanisms such as the 802.11 WEP (Wired Equivalent Privacy) in which an algorithm like RC4 may be used for confidentiality between Mobile Node and "bridge" (access point), but is unlikely to be sufficiently strong both because the key length will be restricted due to export reasons, and there is no provision for a key management protocol with the power of protocols in IPSEC. Ironically our ad hoc system seems to be compatible with the IEEE 802.11b standard.

5.2 "Us" versus "Them"

When one thinks about a secure enclave, one must consider the enclave from two possible points of view. First one must consider mobile systems that belong to "our side". One must then consider mobile systems that may belong to less trusted visitors. Obviously different security policies might exist for wireless (or wired) mobile systems that belong to the home team as opposed to possible potential visitors (or "enemy" crackers trying to gain access via a wireless link available via a passing motor vehicle). For example, one might choose to allow local wireless nodes access to the secure enclave and might disallow visitor nodes. Or one might allow visitors access to the wireless network, but disallow access to key internal secure enclave areas. Certainly any number of policies might exist. It seems that flexibility in this area could be useful. On the other hand, rich security policy implementations could easily be confusing and hard to get correct. We did some policy work here, and probably could have done more (see below under foreign agents for some discussion of one possible interesting security extension). Our work can be summarized in two ways:

1. Our design makes it possible to limit access via "mobile" link-layer interfaces into the secure enclave. Visitors might be totally disabled or allowed access on a case by case basis. This was done via the ad hoc authentication mechanism. A mobile node is required

to know a shared secret (or present a signed beacon) to a agent. If the agent recognized the beacon, it would install a route to give the mobile node access. If the route was not installed, the mobile node might be able to initiate one-way attacks on the enclave, but it would lack the means for two-way communication. Visitors could be given a temporary key to allow them temporary access into the enclave. Our "ad hoc" approach initially requires symmetric keys, but if we were to continue the project for Phase 2, we planned on experimenting with a DNS-based database system for asymmetric keys which would provide more key scalability. The critical idea here is that agents act as routers and do not bridge packets naively into the interior infrastructure.

2. We suggested that network design might be employed to deal with the problem of remote visitors (or local wireless systems) by simply rerouting the internal pipes. For example, any packets coming in on an external unsecured link might simply be tunneled to come back in via an outside external firewall interface. Thus one can easily enable visitors (or local untrusted wireless access) by designing them "outside" the firewall. Packets from trusted external hosts might be allowed direct interior access as long as IPSEC is used (and this risk is deemed worth taking). Mechanisms used here might include known tunnel technologies like CISCO's GRE or IPIP, or even IEEE virtual LANs at the link-layer. The tunnel endpoint (from agent to firewall) must tie to the same sort of input packet filter checks imposed on ordinary Internet packets coming back into the enclave through firewall systems.

5.3 Foreign Agent Considerations

Security policy for foreign agents in our design approach was simple yet very effective. We class foreign agents as either trusted or non-trusted and implemented mechanisms to allow foreign agents to both exclude Mobile Nodes at the link-layer (ad hoc #1) and securely accept tunnel packets from the Home Agent (basically with IPAHIP). Confidentiality was left to the Mobile Node; i.e., the Mobile Node is responsible for making sure that its own packets are secure to/from the Home Agent as it might be using an untrusted Foreign Agent. The reason for using IPSEC authentication in tunnels is to exclude tunnel spoofing possibilities; i.e., the possibility that an attacker might use barebones IPIP to send packets into an infrastructure at an agent, have them uncapsulated, and thus appear to be local with a local IP source address. This is not possible if "our" Foreign Agents only accept IPAHIP packets from Home Agents that they trust and throw away IPIP packets.

We will discuss a more complex security policy system (and implementation) below that we did not choose to implement, but could serve to allow even more complex policies vis-a-vis Mobile Nodes and Foreign Agents.

5.4 Home Agent Considerations

Home Agents serve as a bastion-host for mobile systems. Two-way tunnels terminate (or originate) at the Home Agents. It is assumed that barebones (HA to FA) IPIP tunnels are

not used with Mobile-IP. Instead one ties IPSEC into the tunnel mechanism. We have already discussed how Foreign Agents could insist that all tunneled packets must a priori have some sort of IPSEC association between the Foreign and Home Agent. The Home Agent also serves as the tunnel destination for Mobile Node packets coming back to the enclave. It can enforce a similar semantic; i.e., insist that all inbound IPIP packets must have a Mobile Node/Home Agent (or "our-side" Foreign Agent/Home Agent) security relationship. All "barebone" unsecured IPIP packets would be tossed. Thus the Home Agent can defend the security enclave. One downside here is that Home Agents in this system are not end systems; they are intermediate systems (routers). Thus IPSEC packets may be subject to proposed plaintext attacks, as a "man in the middle" attacker might send packets to the Mobile Node to the Home Agent, and then observe the encrypted packets arriving at the Mobile Node. Defenses against this problem can include session key mechanisms that limit the exposure of keys and/or firewall mechanisms that do not allow Mobile Nodes abroad to talk to systems that are not in the secure enclave.

As a matter of policy, it would be reasonable to assume that Home Agents cannot suffer from a single point of failure scenario. We choose to implement the Home Agent Redundancy Protocol (HARP) so that Home Agents could act in parallel. We did this in such away that IPSEC associations were shared and that in general, Mobile Nodes had no knowledge of HARP.

5.5 Mobile Node Considerations

In summary, we suggest that Mobile Nodes at home might use two-way IPSEC to talk to the Home Agent when an unsecured link is in use. (Note that this implies that a Home Agent should have an exterior and interior secure enclave interface). When abroad, they should use two-way IPSEC tunnels to both defend against malign influences on less secure links, and/or possible interception across the Internet. We regard concerns about "triangle routing" as irrelevant to security concerns. In general, security between parties who have no trust relationship is an oxymoron. The real security policy considerations for systems outside the secure enclave are twofold:

1. Should that system be allowed to talk to home? If the answer, is yes, mechanisms such as two-way IPSEC tunnels could be employed.
2. Should systems that are away be allowed to talk to untrusted systems outside of the secure enclave? If the answer is no, this would obviate such notions as routing redirection targeted to fix "triangle routing".

We also want to point out that our integrated solution includes the possibility of so-called "ad hoc" Mobile Nodes engaged in secure communication. With both our ad hoc systems, Mobile Nodes with a priori trust relationship could setup end-to-end IPSEC tunnels and thus securely communicate using legacy protocols like telnet and ftp. These tunnels were at the network layer, but unlike the Mobile Node to Home Agent relationship, they were end to

end. Thus it is not possible for any attacker to forward packets through a Mobile Node and create a proposed plaintext attack.

Another key idea was the notion in the first ad hoc protocol that an ad hoc network could be based on shared trust. In this case shared trust was obvious as the ad hoc approach uses two shared symmetric keys network-wide. One key was intended for our side and one key was intended to be temporarily created and shared with strangers deemed temporarily trustworthy (and then revoked). However previous digital signature experiment assumed that there was one asymmetric key-pair per Mobile Node. We used a digital signature scheme for all beacons (Mobile Node and Agent), and also used digital signature with Mobile-IP authentication itself. This offers a very novel policy mechanism which, as far as we know, has not been considered before. In our signature scheme, we solved the chicken and egg problem of how Mobile-IP nodes can trust Foreign Agents, by simply asserting that the Foreign Agent's trust statement had to be "mailed" home (sent in the Mobile IP registration) from the Mobile Node to the Home Agent. Thus the Home Agent could use trusted infrastructure to both decide if the FA was trustable and also implement a possible access list control mechanism on non-acceptable Foreign stations. We suggest that the tie between the Mobile Node and Home Agent (as a basic trust duo) is important and can be used to solve many mobility problems. Yes, the Mobile Node is mobile, but it has a fixed surrogate at home that it can interrogate about possibly sticky situations.

Section 6: Suggested Further Work

In this section we are going to present a few ideas that we would like to have pursued but lacked the time and means.

6.1 Mobile Nodes Abroad

It is important to note that Mobile-IP may be viewed as either an Interior Gateway Protocol or Exterior Gateway Protocol (or neither), simply based on security policies. Put another way, it is a reasonable security policy to claim that one is not going to allow foreign Mobile visitors using Mobile-IP (or simpler DHCP loaned addresses) to appear "inside" ones secure enclave. In [2] we noted that anti-spoofing measures currently used in the Internet make cross-domain Mobile-IP problematic. We implemented DHCP-based IPSEC tunnels to show that Mobile Nodes abroad could securely tunnel home and not have any problems with Mobile-IP source address spoofing (the Mobile Node source address is inside the external IPIP encapsulation and will not be seen until the packet arrives home). However we did not implement any mechanisms that would allow an agent to automatically setup tunnels to tunnel non-trusted Mobile Node packets outside the realm of the secure enclave. This is one mechanism for making smarter security agents that could make cross-domain Mobility more feasible.

With such work, those that which to implement this design should pay attention to both verification of the security mechanism and should consider risk assessment for what might happen if such mechanisms are breached. One of the problems with "dynamically poking

holes in firewalls" is that one may get unintended holes. One of the problems with the notion of "active networks" is that they may be more active than one anticipated.

6.2 Smarter Foreign Agents

As another possible mechanism for dealing with both trusted and untrusted Mobile Nodes, Foreign Agents could be made smarter. We suggest two possible optimizations: 1., our ad hoc #1 protocol, served as a very primitive yet very effective mechanism for disallowing cross enclave traffic by Mobile Nodes lacking beacon keys simply because the Foreign Agent would not install a local route for unwanted Mobile Nodes. The end result is that the Mobile Node could send packets but could not receive them through the Foreign Agent. One could improve this mechanism by tying a packet filter access control list to the registration process (802.11 can do this, but at the MAC level). A Mobile Node could present a certificate to a Foreign Agent, and upon verification, the Foreign Agent would then install an ACL that would permit two-way traffic for the Mobile Node. Note that this mechanism should not stand alone (as it is still spoofable) from IPSEC measures. For example, a Foreign Agent might also then insist that all forwarded packets must first be sent directly to the FA itself using a MN/FA IPSEC association. Packets lacking that relationship would be thrown out or unceremoniously tunneled outside the secure enclave. Flexible policy for Foreign Agents may be useful, but appears complex. Of course, the risks inherent in such systems should be considered as well.

6.3 The hard work - integration

To make a long story short, security software is complex and integration although unloved is extremely important. Any key management system seems to have inherent and possibly untoward complexity that informally appears non-linear in nature. This should be recognized and time should be given to security and redundancy oriented research projects that takes these problems into account. Given that we tried to combine Mobile-IP, IPSEC, digital signatures, redundancy in many forms, ad hoc routing, various versions of various operating systems, DNSSEC, etc., it is no wonder that we have had a nightmare integration problem. Our entire software system could stand thorough review and slow and patient replacement of key sub-systems with better quality code. We respectfully suggest that programs oriented towards security and redundancy need to be carefully nurtured and managed in terms of time and budget.

6.4 Keys as a basis for networking

Lastly, our integrated system seems to contain a rather curious notion. In our ad hoc approach, we suggested that IP subnets were not the basis of networks. Instead shared trust (keys ...) should be the basis of networking. Our ad hoc systems could be viewed as

small groups of Mobile Nodes that formed a network based on individual atomic key knowledge of members. Routing in truth may present a chicken and egg problem as you cannot speak ordinary data before you setup routing as control. But routing protocol security is usually solved by an a priori assumption that a set of routers share a shared secret. We suggest it is not unreasonable to assume that all packets might have a trust relationship and that networking might be done from the ground up (even ARP) on that basis.

6.5 Wireless Loading

In the course of the project we discovered that in general wireless LAN did not "load" very well. What we mean is that the number of simultaneous transfers between N laptops and 1 agent is strictly limited and certainly does not scale to the level of ethernet. For example, at the height of our project in terms of group numbers, we encountered a situation with several people present all using a Wireless equipped laptop. All members attempted to simultaneously do an ftp download of a large file. About only half succeeded in simultaneous download. In fact, several of the ftp transfers totally failed, while others would simply wait. While this was an informal experiment we believe it is truly illustrative of the problem. This has never really been a problem for us due to limited numbers of users, few cells, and a narrow distribution of users on campus (located in one corner of a building). However we expect that anyone who widely adopts a wireless LAN link and then faces >5 users in the same cell will experience difficulties when simultaneous bulk transfers are attempted. We expect this problem is due to a number of factors including the limited spreading inherent in current spread spectrum techniques (currently limited by FCC regulations), the fact that the maximum overall bandwidth of 11 MBPS is very distance sensitive, and the fact that such a technology is only capable of "listen while send"; i.e., there is no out of band channel mechanism for true collision detection.

Solutions might lie with techniques borrowed from "Software Radio" techniques; i.e., a very wide spectrum spreading in terms of frequency. More narrowly based devices might use techniques for sampling limited ranges and then switch to another range that does not show so much current use. An agent could easily include the number of current users in its beacon. More research in this area is necessary.

Section 7: Acknowledgements

We would like to thank the following SUNY Institute of Technology graduate students for their participation in the project: Amitabh Pandey, Nikhil Ahluwalia, Shweta Agnihotri and Niranjana Davray. We wish also wish to thank Mr. Anton Gyllenberg of Lifix Go for supplying us with the Beta version Mobile-lp software.

Secure Architecture for Extensible Mobile Internet Transport Systems

Seventh Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 November 2001

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Accomplishments

In this report we present our plans for deploying a combined layer 3 Mobile-IP and IPSEC routing architecture. We discuss possible routing security architectures and then present two alternative designs for an integrated Mobile-IP/IPSEC routing architectures. We refer to these as "closely-coupled" and "loosely-coupled". The closely-coupled architecture relies on a direct binding of IPSEC policy attributes to routing table entries by Mobile-IP routing daemons. The loosely-coupled architecture is based on a more traditional access control list association between the FreeBSD 4.3 Mobile-IP/IPSEC Implementation and a Linux based Mobile-IP/IPSEC integration. Our discussion concludes with an architectural analysis of combined Mobile-IP/IPSEC and a call for the use of IPSEC as part of any mobile VPN scheme.

Introduction

Recently wireless security as found in the popular 802.11 [10] wireless protocol has suffered a series of failures due to the apparent collapse of part of the 802.11 specification called WEP, for "Wired Equivalent Privacy". WEP was intended to offer security services including encryption and authentication. Security experts have recently demonstrated substantial security problems with WEP. Borisov, et. al.[4], describe security and network architecture flaws in WEP. A recent cryptanalytical paper [6] then described a theoretical attack against the RC4 stream cipher used in WEP. Worse, in a recent ATT technical report [17], the theoretical attack was implemented.

The ATT paper in its conclusions suggests that the 802.11 link layer be viewed as insecure. To be fair, the IEEE 802.11 specification stated that WEP was not only optional, but was intended to make wireless "at least as secure as a wire". Unfortunately, this may be misleading to naive users who assume that WEP would offer serious confidentiality services. Borisov's paper, and the ATT technical report both suggest that users should consider using higher-level protocols;

for example, IPSEC [9] and Secure Shell [18]. We concur and further suggest that IPSEC could be directly combined with Mobile-IP [12] in order to make a Virtual Private Network mechanism that is completely based on the layer 3 network layer. This idea was initially proposed by DARPA funded research along those lines at Portland State, between 1995-1999 [15]. In the SAFEMITS project, we intend to explore two different experimental system architectures for a combined Mobile-IP/IPSEC implementation. It should be noted that the earliest architecture was created on the FreeBSD platform using an IPSEC, originally done for NetBSD by the Naval Research Labs, and the PSU Mobile-IP implementation. Our architectural design will be based on the current Linux operating system. We intend to use an existing port of Mobile-IP which was designed to work with Linux to make an integrated IPSEC/Mobile-IP.

Mobile Routing Security Policies

During the first phase of the SAFEMITS project we decided that from a formal point of view, we could distinguish three very general architectural frameworks for mobile routing security. We will call these architectural constructs secure mobile routing architectures. These architectures are as follows:

- MN to security gateway VPN: A two-way VPN is setup between a Mobile Node and some security gateway that acts as an "entry point" into a secure enclave. From the point of view of traditional firewall thinking, the security gateway is a bastion host. In Mobile-IP terms, it may be co-located with the Home Agent (which is the assumption we make in our implementation). Using IPSEC, we setup a two-way layer 3 ESP tunnel, which might or might use dynamic keying. As we will present later in more detail, we have implemented this form of VPN in both of the aforementioned Mobile-IP/IPSEC implementations. Of course, other possible VPN technologies may be used. A Mobile Node outside a secure enclave, has a two-way IPSEC VPN to and from its Home Agent. Foreign Agents are not involved directly in any security association and are merely tunneled over (as are any other layer 3 entities). The first link may be assumed to be wireless, and can be assumed to be outside the secure enclave. The path between the MN and security gateway may be multi-hop and may span the Internet or barring the first link, may be internal to the secure enclave.

In terms of the number and scalability of key associations, key management is linear; that is, for each HA, we have a set of MNs. Key management may be made more complex by security gateway (HA) redundancy issues. We do not rule out a centralized key management system within the secure enclave, that might, for example, use DNS or some other system.

- Agent boundary VPNs: In this form, we restrict cryptographic services to the "external" link; that is, MNs are assumed to be outside the secure enclave, have two-way VPNs between themselves and a boundary agent, and the link connectivity is confined to only one link. The typical boundary agent could have one external link and one internal link. Boundary agents might be layer 3 entities as with Mobile-IP agents, or layer 2 entities as with 802.11/WEP

access points. Boundary agents in a MIP system would insist that MNs must have an a-priori security association. Thus MNs that do not have local IPSEC keys would not be able to penetrate the secure enclave security architecture, FAs, by definition must belong to your security enclave, and MN-FA security associations must exist. Note that in the previous architecture, FAs were not part of the picture. Manual key administration here is fundamentally not scalable. as key associations are a function of the number of boundary agents times the number of Mobile Nodes. We believe that an internal tie-in between IKE daemons and centralized key service, possibly via DNSSEC is mandatory. For example, the Portland State University project did not implement such an architecture, although they did view it as possible future work. However, the PSU researchers did implement a layer 3 authentication system for Mobile-IP itself, that required authenticated ICMP advertisements from all network elements including agents and MNs [2]. Both BBN [19] and SMN also implemented layer 3 authentication systems based on per node digital signatures. Note that such authentication systems are not intended as replacements for higher-order confidentiality systems like IPSEC. They are merely supplemental.

- Secure multi-hop ad hoc routing: Multi-hop ad hoc routing refers to Mobile Nodes that setup multi-hop routing paths via a new class of dynamic routing algorithms; for example, please see [3]. The PSU project implemented a form of DSR [7] in which end host to end host IPSEC associations were manually available. Thus all packets between any two MNs could have IPSEC applied to them. It is important to note that "consenting" MNs in such an architecture, by definition, belong to the same security domain.

In the following section, we present our design of a closely-coupled IPSEC/Mobile-IP architecture. Next we present the alternative loosely-coupled architecture in which we have combined our Mobile-IP with IPSEC. In section 5, we present some architectural analysis in terms of system organization, and finally we present our conclusions.

Closely-Coupled Linux Mobile-IP/IPSEC

We will briefly discuss some architectural aspects of our combined Mobile-IP/IPSEC architecture. Much of the Mobile-IP architecture itself has proved portable over time, but the IPSEC aspects themselves did not survive abandonment of the Naval Research Labs (NRL) IPSEC mechanisms, based on older RFC 1825 IPSEC.

We will briefly describe a new IPSEC mechanism based on close coupling of routes and IPSEC policy. We call this closely-coupled because the route daemons directly manipulate the IPSEC policy. Our design calls for creating an experimental IPSEC policy system based on modification of the route(4) socket.

IPSEC assumes two abstract databases in the operating system, that can be used for cryptographic operations on packets. One may be called a policy database with rules similar to: use IPSEC (tunnel/transport/ESP/AH), on these IP addresses, with a certain security association (algorithms/keys) Formally this database is called the Security Policy

Database or SPD. The other database provides key material; for example, use BLOWFISH, 3DES, with certain key bit strings, and is called the Security Association Database or SAD. Our design for routing socket modifications will allow routes in the routing table to act as the SPD. We assumed key material had a priori been loaded into the SAD. Thus the SPD references the SAD for actual key materials.

Logically the route(8) command could be assumed to have the following form:

```
# route <ipsec-mode> -spi <SPI> -itsrc <SA-ipaddr> -itdst <SA-ipaddr>
```

The ipsec-mode could be any of -ah, -esp, -ahunnel, -esptunnel. The modes defined a particular route as either transport or tunnel mode IPSEC. When a route was loaded, either manually or by a mobile routing daemon, internally a search was performed in the kernel for the SAD, and if an appropriate binding was found, a pointer was setup between the routing table, and the SAD.

The Linux operating system has chosen to adopt the version of IPSEC developed in Helsinki, Finland. This system is based on the RFC 2053 version of IPSEC. Of course, it has also never been burdened with US export law problems.

It should understand that most conventional IPSEC implementations are based on rulesets similar to firewall access control lists. The Mobile-IP/IPSEC routing table feature was used for a number of different security architectures. For example, we plan on implementing the basic MN to security gateway VPN. The HA to MN routing path will require creating an ESP tunnel on the IPIP tunnel device, resulting in packets with an IP ESP (IP datagram) header structure. Other security features will include HA to FA 1-way authenticated tunnels with a IP AH IP datagram structure as opposed to the conventional IPIP tunnel. Also our mobile-node daemon will be capable of using a combined form of DHCP, ESP, and Mobile-IP, when no foreign agents are to be found. This design allows a mobile node to retain its invariant MN IP when away from its home IP address area. The use of the DHCP IP address as the COA meant that any possible IP ingress address problems were avoided because the COA address did not belong to the MN's home addressing domain (see, for example [1], and [5]). Thus, Mobile-IP enabled systems can wander away from their home security enclaves without having to worry about the IP source ingress filter problem.

Loosely-coupled FreeBSD Mobile-IP/IPSEC

The current architecture, based on 4.3 FreeBSD, combines the KAME IPSEC implementation and the PSU Mobile-IP daemons. We have re-implemented the basic MN to security gateway VPN. We refer to this architecture loosely-coupled because the Mobile-IP daemons do not directly manipulate the IPSEC policy. In KAME, IPSEC policy is setup more on the lines of traditional layer 3 access control lists. We assume initial IPSEC two-way tunnels are setup between the Home Agent and Mobile Node, and then run Mobile-IP on top of that configuration. In this section, we will explain the implementation setup in detail, and discuss some resulting implementation problems and solutions.

From the high level point of view, as routing consists of two 1 way problems, we must deal with 2 problems, 1. MN to HA, and 2. HA to MN. For IP datagrams sent from the MN to the HA, we tunnel conventional IP datagrams from the MN to the HA. Thus the IP outer header has an IP src = MN IP, and an IP dst = HA IP. The ESP header encrypts the interior datagram sent from the MN to some other host. For the HA to MN path, we first have packets tunneled via an IPSEC tunnel (IP ESP, IP datagram), where the outer IP header has an IP src = HA IP, and IP dst = MN IP. This packet is then encapsulated inside an IPIP datagram that deals with the COA. Conceptually the HA to MN path can thus be viewed as (IP (dst=COA), IP ESP, IP datagram).

Architectural Details

The FreeBSD 4.3 KAME/IPSEC system allows three levels of kernel control (see `ipsec(4)`). `Sysctl(8)` can be used for global policy. The manual `setkey(8)` command is used to set IPSEC packet-filter defaults { which are similar to traditional layer 3 access control lists implemented in routers. In addition, `setsockopt(2)` can be used for setting per socket IPSEC policy attributes. Thus routing daemons could choose to avoid more general policy when warranted. We make no use of the `sysctl` mechanism and instead use a combination of the `setsockopt(3)` and `setkey(8)` mechanisms.

Mobile-IP interoperation

The basic MN to HA two-way VPN policy requires several modifications to the Mobile-IP daemon implementation. First of all, as one possible security policy choice, we chose to make the necessary UDP and ICMP sockets, bypass any and all IPSEC packet mechanisms in the kernel. This is done using `ipsec set policy(3)`, and `setsockopt(2)` calls. This means that all Mobile-IP packets bypass local IPSEC, and must rely upon their own devices for security. Remember that we choose to ignore Foreign Agents, thus it is important that we be able to talk to them and not assume we speak IPSEC with them. Further, by definition, we cannot share secrets with agents from another security domain. Hence we choose to let Mobile-IP as a protocol stand on its own, otherwise MNs would wrap Mobile-IP registration packets in ESP, FAs might not understand them, and thus could not relay them to the Home Agent.

The second implementation aspect is unfortunately far trickier. When a MN visits a FA, "all" packets in theory will be delivered via a HA tunnel encapsulation; that is, datagrams processed by the KAME IPSEC tunnel are formed as IP ESP f IP datagram g, with the outer IP src = MN IP, and the outer IP dst = HA IP. Unfortunately this runs full tilt into the BSD ARP table implementation. In the current BSD architecture, the ARP table is not separate from the routing table, and is implemented via a so-called clone route mechanism. When an interface uses ARP, and its IP address is configured, a clone route is placed in the routing table. For example, in the sample routing table below, we see a clone route (marked with the UC tags) that was loaded for a local ethernet

interface when the interface was booted. One ARP table entry was instantiated in the routing table for local IP address 10.0.0.1, and later filled in by the ARP protocol itself with the MAC address of the local link host, 10.0.0.1.

```
host# netstat -rn
Destination Gateway Flags Netif Expire
10.0.0.0/8 link#1 UC 0 0
10.0.0.1 0:d0:c0:5b:18:0 UHLW 4 3
```

This means that when a MN visits a Foreign Agent, and the first packet is sent via the IPSEC ESP tunnel from MN to HA, the outer IP header will of course, have IP dst = HA IP. This in turn, will cause the clone route to create an ARP table for the HA, because the MN after all, shares a local IP subnet association with the HA. Naturally since the HA is not nearby, this causes complete failure as no packets can reach the HA.

In order to fix this problem, we will need to modify mnd to take advantage of the state machine. When configured for IPSEC, and in NOWHERE or AWAY states, it simply deletes all ARP table entries, and also deletes the clone route. When at home, the clone route is reinstalled. This is one possible policy choice, and the implementation might eventually allow more flexible configuration policies.

Architectural Analysis

In this section we wish to present an architectural analysis and briefly consider two questions:

1. what key ideas might be necessary in an operating system architecture to allow a combined Mobile-IP and IPSEC?
2. What are the pros and cons of the two Mobile-IP+IPSEC architectural approaches, themselves?

We suggest that KAME IPSEC has provided us with two necessary features that we hope would be available with any IPSEC implementation. The first feature that was important is the ability to specify with the KAME packet filter mechanism that "all packets" should be sent over a tunnel to a tunnel endpoint. For example, the MN should be able to send "all" packets to the HA. It is hard to imagine that an IPSEC implementation would not have this capability, but it is fundamental and necessary.

The second important IPSEC capability is the ability to override higher-level "all packets must use IPSEC" packet filters on a per-socket basis. Without it, Mobile-IP registration packets could not be relayed by Foreign Agents that do not belong to the security domain. More generally, it is extremely reasonable for routing daemons using any routing protocol to be able to except themselves from a system-wide IPSEC policy. Most routing protocols have their own authentication mechanisms; for example, OSPF [11] has per link authentication.

We have arrived at the conclusion based on our preliminary analysis that the "Helsinki" Linux based IPSEC/MIP experimental implementation was strongly-coupled, because the IPSEC policy was directly manipulated by the mobile routing daemons. On the other hand, the FreeBSD 4.3 KAME/IPSEC MIP is loosely-coupled. KAME IPSEC handles most of the IPSEC-based tunneling. We assume that the KAME IPSEC has been setup, and then run Mobile-IP which merely overrides any IPSEC policy in order to get Mobile-IP functions themselves accomplished.

So the bottom-line question then remains: which is better? It has been said in the past that any "packet filter" or access list mechanism vis-à-vis firewalls may be dangerous, because if the rule set is complex, it is easy to make mistakes. The FreeBSD 4.3 KAME/IPSEC route-based mechanism however is perhaps more esoteric than any possible ACL mechanism. On the other hand, the bottom-line issue here may simply be portability. Our Mobile-IP implementation when ported to Linux will make a very few, reasonable assumptions about IP mobility features needed by an operating system. By comparison the PSU FreeBSD Mobile-IP/IPSEC implementation did not lend itself to simplicity or portability as it made complex assumptions about the host OS IPSEC and routing socket architectures. Our design is more elegant and much simpler.

Summary

In a narrow sense, we do not know of any related work, other than the comparison of the FreeBSD 4.3 version to the Linux OS version as previously presented. In a wider sense, we could consider competing link-layer, and network-layer systems that are somehow targeted at mobile security

Such systems could include 802.11 WEP (link-layer), or other VPN systems like PPTP, which has been fairly well discredited by Bruce Schneier [14]. The critical question is this: Why is IPSEC not chosen by default as the main vehicle for the delivery of end system to border gateway virtual private networks?

IPSEC has major virtues including:

1. It is not specific to any link-layer, and could be used for cellular telephony wireless, or over Ethernet for that matter.
2. It is not link-specific in terms of hop counts. It can easily be multi-hop across the Internet to a remote home security enclave.
3. It has been widely and opened reviewed in the IETF.

4. Over time, it will improve or at least keep up as it was designed for both replacement of its basic cryptographic algorithms, and key exchange algorithms. Thus it is more extensible than fragile algorithms like WEP.

It may be argued that from the layer 1 and layer 2 "IEEE points of view", the IEEE cannot assume that IETF protocols are in use. What would be wrong then with doing nothing? KISS has its virtues and trying to put complex functions like security into firmware or hardware may be best left to layers 3 and above. One might argue that combined IPSEC/Mobile-IP is not a good combination, because perhaps Mobile-IP is not a good idea. There are those who argue that Mobile-IP may perhaps be inefficient or have other problems. It is not our goal here to argue for or against Mobile-IP. One could just as well combine IPSEC with DHCP. DHCP could be authenticated itself, or perhaps protected by IPSEC in local security domains, and then IPSEC could take care of the two-way tunnels to and from a home security agent. Obviously with DHCP, and unlike with Mobile-IP, IPSEC cannot take advantage of a fixed IP address as a index mechanism because DHCP IP addresses may vary over time or over link reattachments. However, IPSEC provides for this possibility with its dynamic key management protocol called IKE. According to the IPSEC Domain of Interpretation [13], one can simply setup two-way tunnels with IPSEC using dynamic keying and a fixed higher level name a la "user@dnsname", or according to the DOI document, ID USER FQDN. Again, there is no point in avoiding IPSEC.

References

- [1] J. Binkley, and J. Richardson, Security Considerations for Mobility and Firewalls, IETF draft, 1998,
<http://www.cs.pdx.edu/jrb/jrb.papers/firewall/draft.txt>.
- [2] J. Binkley, and W. Trost, Authenticated Ad Hoc Routing at the Link Layer for Mobile Systems, Wireless Networks, Vol. 7, No. 2, pp. 139-145, 2001.
- [3] J. Broch, D.A. Maltz, D.B. Johnson, Y.C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols", Mobicom, Dallas, October 1998, pp. 85-97
- [4] N. Borisov, I. Goldberg, D. Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11", In Proceedings of MobiCom 2001, July 2001.
- [5] P. Ferguson, and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC 2267, IETF, January 1998.
- [6] S. Fluhrer, I. Mantin, A. Shamir, "Weaknesses in the key scheduling algorithm of RC4". Eighth Annual Workshop on Selected Areas in Cryptography, August 2001.
- [7] David B. Johnson and David A. Maltz. \Dynamic source routing in ad hoc wireless networks", In Tomasz Imielinski and Henry F. Korth, editors, Mobile Computing, pages 153{181. Kluwer Academic Publishing, 1996.

- [8] KAME IPv6 and IPSEC project,
<http://www.kame.net> , Sept. 21, 2001.
- [9] S. Kent, and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, IETF, November 1998.
- [10] Local and Metropolitan Area Networks, IEEE. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Standard 802.11, 1999.
- [11] J. Moy, "OSPF Version 2", RFC 2328, IETF, 1998.
- [12] C. Perkins, "IP Mobility Support", RFC 2002, IETF, 1996.
- [13] D. Piper, "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, IETF, 1998.
- [14] B. Schneier, and Mudge. "Cryptanalysis of Microsoft's Point-to-Point Tunneling Protocol (PPTP)", 5th ACM Conference on Computer and Communications Security, pp. 132-140, San Francisco, California, November 1998. ACM Press.
- [15] Secure Mobile Networks project,
<http://www.cs.pdx.edu/research/SMN> , Sept. 21, 2001.
- [16] K. Sklower, "A Tree-Based Packet Routing Table for Berkeley UNIX", Proceedings of the 1991 Winter USENIX Technical Conference, January 1991.
- [17] A. Stubblefield, J. Ioannidis, A. Rubin, "Using the Fluhrer, Mantin, and Shamir Attack to Break WEP", ATT Labs Technical Report, TD-4ZCPZZ, Revision 2, August 21, 2001.
- [18] T. Ylonen, "SSH - Secure Login Connections over the Internet", USENIX Security Conference VI, 1996, pp. 37-42.
- [19] J. Zao, S. Kent, J. Gahm, G. Troxel, M. Condell, P. Helinek, N. Yuan, and I. Castineya. A Public-Key Based Secure Mobile-IP. MobiCom97, September 1997.

APPENDIX B

SUNY Institute of Technology at Utica/Rome

Secure Architecture for Extensible Mobile Internet Transport Systems

Tenth Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 February 2002

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Project Status Overview

1.1 Redundancy

We are going to briefly touch on the current theory behind our redundancy work. We have "theory" in place for the following areas:

1. multi-hop routing
2. home agent redundancy
3. foreign agent redundancy.

In addition, we have realized that redundancy could also focus on allowing a Mobile Node to bind more than one interface at a time to a route (especially the default route). We do not intend to currently pursue that area, but the notion is interesting and could be useful.

For now however we only wish to raise a few points:

1. We would consider using a source-route based protocol that will be based on top of our current ad hoc protocol. The current protocol is authenticated and establishes link-layer reachability. It is subnet-free. The next generation will also be authenticated and subnet-free. We hope to be able to explore multi-path aspects; i.e., give this protocol a redundancy orientation.

2. We have encountered a difficult kernel routing implementation problem. Basically we need an "input" (besides an ICMP host unreachable message, which may not arrive) from the routing code in the network layer in order to drive the "on demand" source route query in the MN routing daemon. The problem is that if a default route is present, it always matches any destination lookup where there is not a more precise answer (a specific host route). If the default route is not present, the kernel will generate a "route miss" message and send it upstream to readers of the route socket (routing daemons). It is not satisfactory to simply remove the default route, since for MNs that can hear an agent, the default route will be set to the agent.

We have experimented with adding a function to the routing code of the kernel that will cause the kernel to generate a "route clone" message whenever cloning occurs. This will allow us to learn that the kernel has generated a "host route" when the default route is used, and we can use this as an input to drive the multicast source query mechanism at a Mobile Node.

1.2. Home Agent Redundancy

The HA redundancy specification work is started. For a change, we do not need any kernel work. We think that HA redundancy might take two approaches:

1. "serial" HA redundancy
2. "parallel" HA redundancy.

"Serial" means that a MN would have a list of two HAs and would only use one HA at a time, but would be able to switch between them, if the current HA failed to respond to a Mobile IP UDP request. Parallel would mean the use of two HAs at once. We prefer the former, since it would lighten the Internet load and seems more practical. The basis of the idea here is that the HAs will use normal unicast IP routing and act as unicast routers to the MN subnet, which will be partitioned by definition (i.e., the wireless subnets on the other side of the HAs will not be co-located). In other words, we intend to use normal unicast IP routing for reachability to the Mobile IP subnet. The important point here is that at any time, an individual packet from a host sent to a MN might go to either one of the HAs. As a result, the HAs must keep each other up to date about the location of MNs.

We intend to configure the two cooperating HAs so that know about each other and keep a TCP connection open between the HA pairs. They will use that connection to exchange information about MN status, as they must keep a parallel routing table setup for all cooperating MNs. A TCP-based application routing protocol will be defined that allows the HAs to check that their peer HA is up, and exchange MN location information. The HAs will install tunnels in parallel for MNs at remote FAs. If an MN is actually "home" at one HA, the peer HA will install a tunnel to the other HA for packets that are naturally routed to it. One idea that we should consider here is whether or not MNs could be dynamically informed of a new HA. This would allow for a smooth handoff in case one HA was taken down and replaced by another that did not share the same IP address for the HA. This is a feature that could be built on top of the "serial HA redundancy" system. On the other hand, this may be a feature that is not useful in the short term.

1.3. Foreign Agent Redundancy

We have already identified one possible aspect of FA redundancy that has already been implemented and available, a "faster" handoff algorithm that allows us to switch between FAs and gives us some ability to deal with FAs that are partitioned from our current HA. We hope to soon (probably in spring) begin an experimental implementation that will allow a MN to talk to two FAs at the same time. The fundamental feature point needed here is that we must teach the routing code in the kernel to allow us to load a second gateway. Logically we want to be

able to bind a list of gateways to a given destination. IP output would send each packet to every gateway in the list.

In terms of the routing table, we might have:

Destination IP	Gateway 1/if1	Gateway 2/if2
1.2.3.4	2.2.2.2/ed0	3.3.3.3/wlp0

When IP sends a packet to destination 1.2.3.4, we would send two packets. One would go the next hop router at 2.2.2.2 via the Ethernet ed0 interface. The other would go the next hop router at 3.3.3.3 via the wireless interface. (We would probably use the same interface in both cases, but the capability to use two interfaces is interesting). We are simply trying to do what we can to improve the odds that a given packet might make it. On the other hand, the cost to the network is obviously twice as much. Still this may be useful in extreme cases, and as is often the case, a little mechanism in the kernel may provide a lot of flexibility at the application layer. We intend to test the mechanism in the next couple of months and think about how to use it in the long term.

FA redundancy would use this mechanism as follows:

1. at the HA, the HA could setup a tunnel route to the MN, and bind two FA gateway/COAs to it.
2. at the MN, a MN using Wireless LAN currently has a list of agents (FAs) sorted by signal strength. Instead of using the top single agent, we would use the top two.

We expect this dual gateway mechanism would make intermittent connectivity "better". Possibly it should only be used in the case of intermittent connectivity to the top two FAs. How to measure any cost and advantage here remains an open question.

1.4. Rough Measurements of IPSEC and Mobile-IP over Wireless LAN

We have at initial cut at implementing an IPSEC/Mobile-IP integration on a few Pentium machines, over Wireless LAN, and over "localhost". "Localhost" means we run the client and server on the same machine, thus we get a rough idea of a machine's compute power. The goal here was simply to get a sense about IPSEC in terms of algorithms and to get a feeling about the cost of IPSEC encryption over the Wireless LAN. These are rough benchmarks and it is important to point out that the conclusions are rough as well. We will first characterize the test environment in terms of hardware and software tools, then present the numbers, and finish with a few conclusions.

Hardware/hosts:

1. SONY Vaio laptops, 750 mhz Pentium
2. Home agent running on 233mhz Pentium machine.

Hardware/Wireless 802.11LAN:

In theory, 802.11 Wireless LAN is advertised as having an 11 Mbps speed. We typically see speeds between 4 Mbps and 10 Mbps depending on collisions and on timeout errors that are built-in to the LAN controller used in the SONY wireless system. The testing here was done when the local Wireless LAN link was otherwise unused. We were able to verify this by using a real-time network analysis tool that shows local link traffic. As a specific example, we typically see speeds of 100k bytes per second with ftp transfers over Wireless LAN.

Software/IPSEC route binding:

Using the arp command, we installed routes between the two hosts over the Wireless LAN links. We used ftp to transfer data.

Software/tcpclient + tcpserver:

As part of our IPSEC verification effort, we created a tcpclient/tcpserver application pair that allow us to send various sized tcp packets over a TCP connection. We used this to test the Mob/IPSEC transport mechanism. Tcpclient sends packets to tcpserver that discards them.

Software/ssh:

Ssh is a tool that allows RSA-based authentication, followed by use of encryption across a TCP-based connection. By default, ssh uses idea with 128 bit keys for encryption but can be told to use DES or 3DES as well. Ssh is entirely an application layer application. It does not use IPSEC and might well be considered as a competitor for IPSEC.

We have not yet benchmarked IPSEC/ESP/DES over Ethernet and hope to be able to do that this coming month. Our primary hypothesis is that the computationally intensive cryptographic algorithms DES and 3DES are still fast enough on a Pentium machine so that Wireless LAN does not become the bottleneck. It is possible that faster Wireless LAN systems may come onto the market in the next few years, but most likely compute speed will go up as well, and probably will be able to keep up with the encryption cost in the near future.

On the other hand, a Home Agent acting as the tunnel endpoint for N remote Mobile Nodes using DES/tunnels back to the HA, will need all the compute power it can get. We will attempt in the next quarter or two to measure simultaneous DES stress on the Home Agent (if we can figure out a meaningful test case).

2.0 Suggested Further Work:

In this section we are going to present a few ideas that we would like to have pursued but lacked the means or time., however, would be strong considerations if the SAFEMITS project is funded for an additional year.

2.1. Mobile Nodes Abroad

It is important to note that Mobile-IP may be viewed as either an Interior Gateway Protocol or Exterior Gateway Protocol (or neither), simply based on security policies. Put another way, it is a reasonable security policy to claim that one is not going to allow foreign Mobile visitors using Mobile-IP (or simpler DHCP loaned addresses) to appear "inside" ones secure enclave. We noted that anti-spoofing measures currently used in the Internet make cross-domain Mobile-IP problematic. We implemented DHCP-based IPSEC tunnels to show that Mobile Nodes abroad could securely tunnel home and not have any problems with Mobile-IP source address spoofing (the Mobile Node source address is inside the external IP encapsulation and will not be seen until the packet arrives home). However we did not implement any mechanisms that would allow an agent to automatically setup tunnels to tunnel non-trusted Mobile Node packets outside the realm of the secure enclave.

This is one mechanism for making smarter security agents that could make cross-domain mobility more feasible. With such work, the implementers should pay attention to both verification of the security mechanism and should consider risk assessment for what might happen if such mechanisms are breached. One of the problems with dynamically poking holes in "firewalls" is that one may get unintended holes. One of the problems with the notion of "active networks" is that they may be more active than anticipated.

2.2 Smarter Foreign Agents

As another possible mechanism for dealing with both trusted and untrusted Mobile Nodes, Foreign Agents could be made smarter. We suggest two possible optimizations:

1. our approach served as a very primitive mechanism for disallowing cross enclave traffic by Mobile Nodes lacking beacon keys simply because the Foreign Agent would not install a local route for unwanted Mobile Nodes. The result (again) is that the Mobile Node could send packets but could not receive them through the Foreign Agent. One could improve this mechanism by tying a packet filter access control list to the registration process (802.11 can do this, but at the MAC level). A Mobile Node could present a certificate to a Foreign Agent, and upon verification, the Foreign Agent would then install an ACL that would permit two-way traffic for the Mobile Node. Note that this mechanism should not stand alone (as it is still "spoofable") from IPSEC measures. For example, a Foreign Agent might also then insist that all forwarded packets must first be sent directly to the FA itself using a MN/FA IPSEC association. Packets lacking that relationship would be thrown out or unceremoniously tunneled outside the secure enclave. Flexible policy for Foreign Agents

may be useful but appears complex. Of course the risks inherent in such systems need to be considered as well.

2.3. The hard work - integration

To make a long story short, security software is complex and integration although unloved is extremely important. Any key management system seems to have inherent and possibly untoward complexity that informally appears non-linear in nature. This should be recognized and time should be given to security and redundancy oriented research projects that takes these problems into account. Given that we tried to combine Mobile-IP and IPSEC using various versions of various operating systems, FreeBSD, Linux, Windows, etc., it is no wonder that we have had a nightmare integration problem. Our entire system could stand thorough review and slow and patient replacement of key sub-systems with better quality code. We respectfully suggest that programs oriented towards security and redundancy need to be carefully nurtured and managed in terms of time and budget.

4. Keys as a basis for networking

Our design approach seems to contain a rather curious notion in that we suggested that IP subnets were not the basis of networks. Instead shared trust keys should be the basis of networking. Our test bed system could be viewed as small groups of Mobile Nodes that formed a network based on individual knowledge of members. Routing in truth may present a "chicken and egg" problem as you cannot transmit ordinary data before you setup routing as control. But routing protocol security is usually solved by an a priori assumption that a set of routers share a shared secret. We suggest it is not unreasonable to assume that all packets might have a trust relationship and that networking might be done from the ground up (even ARP) on that basis.

2.5. Wireless Loading

In the course of the project we discovered that in general wireless drivers did not "load" very well. What we mean is that the number of simultaneous transfers between N laptops and 1 agent is strictly limited and certainly does not scale to the level of ethernet. For example, if all project team members attempted to simultaneously do an ftp download of a large file, only a few succeeded in simultaneous download. In fact, several of the ftp transfers totally failed, while others would simply wait. This was an informal experiment but we believe it is truly illustrative of the problem. This has never really been a problem for us due to limited numbers of users, lack of cells and a relatively confined number of users on campus (in one building and one lab). However we expect that anyone who widely adopts a wireless LAN link and then has > 5 users in the same cell will experience difficulties when simultaneous bulk transfers are attempted.

We expect this problem is due to a number of factors including the limited spreading inherent in current spread spectrum techniques (currently limited by FCC regulations), the smaller overall bandwidth (≤ 11 Mbps), and the fact that such technology is only capable of "listen while send"; i.e., there is no out of band channel mechanism for true collision detection.

Solutions might lie with techniques borrowed from DARPA "Software Radio" techniques; i.e., a very wide spectrum spreading in terms of frequency. More narrowly based devices might use techniques for sampling limited ranges and then switch to another range that does not show so much current use. An agent could easily include the number of current users in its beacon. More research in this area is necessary.

APPENDIX C

SUNY Institute of Technology at Utica/Rome

Secure Architecture for Extensible Mobile Internet Transport Systems

Ninth Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 January 2002

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Project Status Overview

IPSEC, Mobile-IP and Policy

We wish to preface our remarks by pointing out that it would be useful from a technology transfer perspective to confer with industry, firewall vendors, and IPSEC software developers, about what kind of policy would make sense in terms of Mobile-IP. Once we know, we can try and implement it. Up until now, we have addressed a very limited range of topologies, but as we begin to explore "dynamic key" issues, it would be most useful to try and determine what actually might be useful and/or feasible. We shall attempt a short discussion here in order to begin to explore some of the issues. We recognize that this discussion will not be complete, but it is a start. Hopefully it might fuel some high level

As a word of caution, the issues for security and Mobile-IP are complex and

are also subject to change. Mobile-IP itself is a moving target as recently the subject of a "reverse tunnel" protocol setup by MNs where the local FA tunnels back to the HA has come up in the IETF Mobile-IP working group. As is often the case, it seems little attention is paid to security issues, even though an entirely new topological dimension has been revealed.

For the sake of argument, let us first establish some topological notions. We will try to bind our policies according to the topologies. In conclusion, we will try to summarize any lessons learned. Mobile-IP topologies of interest might include:

1. MNs/HAs.

.at home

---- on a wireless system (Wireless LAN)

---- on a wired system (ethernet)

.away

-----via a FA

-----via a COA (but no FA)

2. MNs/FAs. It is possible that a FA might have a pool of COAs. This might be useful although Mobile-IP currently does not address it as an possibility.

3. MNs/COA. Here an MN acquires a COA as a local link IP address. The IP address may be acquired via DHCP or some other mechanism including PPP and manual administration.

4. FAs/HAs

5. ad hoc MNs

6. route optimization (MNs talking directly to CHs).

There are two kinds of protocol considerations that are fundamental. One can fundamentally distinguish between "control" packets (routing including link-layer mechanisms like ARP), and "data" packets that can be transmitted after routing is established. Routing must be setup first. In certain topologies (barring ad hoc), this is what Mobile-IP attempts to do. This is of course, the chicken and egg problem that has referred to elsewhere. Without routing, data cannot be sent.

In the discussion below, we will mostly neglect any Policy that says "you may not do this," and always neglect Policies of the form "we don't care if you do this."

It should also be noted that we have intentionally disregarded the useful area of IPSEC/Firewall policy combinations since it is outside the scope of the current project.

1 MNs/HAs

1.1 At home policies:

1. A wireless link may be felt to be less secure than a wired link. As a result, one might require wireless security along the usual lines with the usual possibilities: encryption, authentication, session-keys.
2. Ideally such security should encompass both routing and normal data.
3. For reasons of scalability, one may very well require dynamic key management (as opposed to static). For example, if all employees for the Federal Government have wireless laptops, that will be a lot of keys. (There is not much difference between /etc/keys and /etc/hosts from the scalability point of view.)
4. One may view a Mobile system as a threat simply because it can get up and walk away, possibly with valuable information. Hence it may be subject to any number of special restrictions even when at home (e.g., can only use subnet X, and must use Mobile-IP authentication to even get routed).
5. One may have policies for the HA as a router that limits ingress to interior non-mobile networks.

Observations: Each relationship (like any other network conversation) is symmetrical. Here we have MN to HA and HA to MN. As perviously pointed out, in some cases, one of the sub-relationships may be less important. In this case, the HA is probably more important simply because it represents an entrance to the secure enclave. (This is of course, even more important when we are concerned with routing to and from the HA when an MN is abroad).

One might view the MN at home as being within the secure enclave, unless one is concerned about uncontrolled wireless exposure or uncontrolled taps into local links. The MN at home is not very different from conventional systems at home. ARP is a security exposure anywhere, although it may be deemed adequate for links like ethernet. Certainly small-scale static key distribution for both control/data is possible. But if dynamic key distribution will occur anywhere, it certainly can occur here.

Away Policies:

1. Mobile-IP may be banned. Of course, such a policy may be viewed as not interesting but it is of fundamental importance. Until recently firewalls existed to keep spoofers out and one form of spoofing is borrowing IP addresses that should be "inside." Therefore Mobile-IP can be viewed as institutionally (IETF) approved spoofing. Certainly many organizations will never permit it either for their employees to get back in or for "visitors" from abroad. We will not refer to such a negative policy again, but it is a fundamental challenge for IPSEC to see what inroads if any can (or should) be made about such world views.

2. One may require the MN to authenticate its control packets to the HA (and vice versa).
3. One may require the MN to secure all of its packets including data and either:
 - (a) send packets to home systems securely but not care about packets sent to CH systems not at home
 - (b) Send all packets home first so that the MN appears to originate to CHs from home. The enclave itself may choose to not allow CHs to forward packets to MNs to suppress proposed plaintext attacks through the HA.
 - (c) Only talk to systems at home.
 - (d) never talk to systems at home.

Note that again "secure" can be extended to include authentication, encryption, and session keys.

4. One may prohibit the MN form any sort of security relationship other than with Home. For example, MNs may use FAs but may not form SAs with them, and must not have the FA forward packets into an enclave (reverse tunnel). This is because we choose to not trust non-local FAs.
5. One may require the MN to hide any IP addresses that are native to "home" since exposure of such address topology may lead to external attacks.

Observations: We are talking about the MN and the HA when the MN is not present on the HA's subnet. The MN/HA relationship is potentially useful in many ways, simply because the MN can pre-establish a security relationship with the HA before it wanders abroad. We can also view a MN that is abroad as creating a possible extension of a local secure enclave. Another variation on this theme is that the MN and HA by definition belongs to one administrative/security domain. As we have pointed out previously, ARP spoofing could be viewed alone as a denial of service attack, but if coupled with a previous acquisition of any home IP address or telnet/FTP password sent in the clear, it may become a very potent means of attack on the home enclave.

As a side issue, note that not giving out IP addresses for home systems is an important attribute as well. IPSEC as a mechanism here is very important. We also need to be able to dynamically setup "webs of trust" that include HA-FA-MN pairings.

Specification of routing policy based on IP addresses that leads to exible routing could be useful. One might want to talk to CHs directly and avoid home or send everything home, etc.

1.2 MNs/FAs policies:

1. One may rightfully worry about wireless links again, in that they may be deemed less secure. If a network security officer cannot control wireless links at home, he can certainly not control foreign links. We may hence require link security when abroad and might also ban ARP. Foreign links can be viewed as much more dangerous.
2. A reasonable rationale for the existence of FAs is to charge MNs as customers. As a result, a likely policy here is that any MN will have to use a very good cryptographic means to prove its identity. Dynamic negotiation of trust (for identity as well as usage) would be a requirement.
3. Local security may require that FAs form a closed trust system (no outside MNs may use them).
4. Local security may require that FAs be outside of any secure enclave (and hence there may be no need for security relationships).
5. Certainly the home security organization might desire that MNs reveal as little as possible about themselves to a FA. This might include a requirement that the MN is not to form a SA with a FA.
6. FAs "detunnel" IPIP packets. One may require a SA between the tunnel endpoints.

Observations: It is not clear in how many ways an MN and FA might combine (or not) in security terms. It is clear that it runs the gamut from complete and dynamic security both for routing control and data (wireless link) to nothing at all. In our MobileIP/IPSEC architecture it is possible to send all packets to home with IPSEC and simply use a FA as a router. We could currently setup our FAs to form a "closed" trust system since we have implemented MN/FA static keys. By definition, this would exclude MNs from other domains. However we currently have no means to dynamically negotiate anything between MNs and FAs.

Note that the relationship could be asymmetrical. The MN might not care so much about whether or not the FA can prove itself, since the MN is very much at the mercy of the FA's. If the FA will ship packets out and in, that may do from the MN's point of view. The FA however might care greatly since it may represent a portal to a "secure enclave" or might be interested in economic transactions.

On the other hand, the truly paranoid MN may be concerned about the FA since the FA is certainly a "man in the middle" and might do unpredictable things to the MN's packets. The MN maybe able to take advantage of its permanent trust relationship with the

HA and use the HA to determine the FA's trustworthiness. Is "ABC FA" really a FA? Possibly such a determination could be used to setup two-way trust between both the MN-FA and FA-HA.

1.3 MNs/COAs policies:

1. An MN might be required to ONLY use a COA IP address on outer IP packets sent home so that it can tunnel packets through a border router. Thus even the MN's IP address could be hidden from view.

2. The MN may be required to setup a SA based on the COA with the HA. This would prevent external unsecured IP addresses from entering a secure enclave. Such a policy could extend to both control and data.

Observations: An MN may acquire a COA via some TBD mechanism (DHCP). It might even do so at a FA. Some of the previous Policies apply here (so we left them out). There is however a new wrinkle, and that is that the MN has acquired an IP address with local significance. This is a double edged sword in that it may be useful in that an MN with a local COA can use it to bypass an ingress filter but one could view such a COA mechanism as a security liability.

1.4 HAs/FAs policies:

1. One may require either a static or dynamic trust negotiation for at least routing packets (MIP). Dynamic keying would be needed for FAs not under local control.

2. IPIP packets may not be permitted. One might require IP/IPSEC-IP or IPIP/IPSEC where there is some way to bind at least one of the local headers to the box "at home" that performs the IPSEC operation.

Observations: We feel the most important point to make here is that IPIP can be viewed as a liability especially across security domains as it is basically a mechanism for directing an IP datagram anywhere (and shedding the IP outer header or skin that got the datagram to the tunnel endpoint). One could use conventional firewalls to protect internal "detunneling" devices that lack protection of their own. On the other hand, IPSEC is an obvious mechanism for making IPIP safer. Further one needs dynamic key distribution to enable IP IPSEC IP as a mechanism for boxes not in the same security domains.

1.5 ad hoc MNs policies:

1. MNs should authenticate routing packets.

2. MNs should secure control information.

3. MNs should only talk to their own kind and maybe only when they see their own kind.

Observations: The first two policies above are nothing new and other policies previously mentioned can certainly apply (e.g., concern about wireless links abroad). It is important to note that MNs in an ad hoc situation may be worse off than an MN at a FA, simply because they can be cutoff from the wired infrastructure. Both DNS and any existing Certificate Authority may not be available. (The lack of DNS certainly stops some not very bright applications from even running).

1.6 Route Optimization policies:

1. We may simply disallow it as we do not want our MNs to talk to CHs except through home. Of course we may not want them to talk to non-home CHs either.

2. Control packets should be authenticated.

3. Previous concerns about tunnels apply as well.

Observations: One can point out that routing optimization is an interesting name for a noble attempt to make MNs behave like ordinary hosts; i.e., avoid the routing triangle problem. However from a security point of view, one might not only like CH packets to go through the HA (you could log who the MNs are talking to), but worse you may want all of your traffic from the MNs to come back home first as well. Hence one would have a double triangle (MN - HA - CH and CH - HA - MN) and be satisfied with it.

Routing optimization is a form of source routing (as is any use of IPIP tunnels) and one might think that minimally authentication for tunneled packets is a requirement.

1.7 Summary

1. MNs in ad hoc situations and at FAs are at a disadvantage because they either lack connectivity or must do it through a man in the middle. How do you get a certificate from the DNS if there is no DNS? You may have to bring something with you to either minimize exchanges due to lack of bandwidth or fast motion, or you may simply not have access to the wired infrastructure in an ad hoc situation or you may have to establish a "face-to-face" web of trust.

2. MNs can take advantage of the a priori security relationship with the HA. Our position is that a MN need not ship a certificate to the HA to prove itself since it can do something more lightweight and/or prepare a bit before it leaves. The MN can also ask the HA about a FA since the M-IP protocol itself is extensible. Some sort of "in-line keying" could be done there in the MIP protocol. A FA might ask a HA about an MN (as well as about the HA itself) as we can presume that the FA/HA bandwidth might be greater than the MN/FA bandwidth.

3. Tunnels need to be secured by IPSEC and be setup via some dynamic exchange of key material since security domains may be crossed. Note that Home Agents may "detunnel" packets as well as Foreign Agents.

4. If policy dictates that FAs need keys, FAs need more than manual key mechanisms, although such a mechanism may have limited utility in a small network infrastructure. FAs need to dynamically negotiate with HAs in order to setup tunnels and to make sure that all packets coming through a tunnel are secure. This is an absolute must. FAs may need to negotiate with MNs where a security policy dictates that MNs must authenticate themselves and must negotiate with MNs wherever FAs might charge for their services.

5. One must consider "scalability" in terms of manual keying. This may be deemed obvious, but one will have more FAs than HAs, and if you use manual keying with FAs, you must add each new FA key to all MNs and teach that FA about all MNs. This is less scalable than the MN-HA relationship. Manual key distribution may also be hard to automate. It can be done where FAs are concerned since we assume they are reachable. But given that laptops may be off or away, the laptop side is very hard to automate in practical terms.

6. In an environment where an MN acquires and uses a COA, the neighborhood local router might choose to impose security policies on the MN. The kinds of policies between such a router and a FA may not differ very much.

7. This last set of router policy considerations may be extended to filtering border routers and might require some TBD protocol that can discover the number of policy-oriented routers between here and there (and back). Such a system begins to sound similar to other problems facing the Internet (e.g., RSVP setup management, or nested tunnels and "soft state"). We need to consider the dangers of creating a system that requires "hard state" end to end, so that if dynamic routing must occur, a "connection" will be lost. Such a system would not be very much like the current Internet and might be very hard if not impossible to debug.

Future work

In this section we are going to present a few ideas that we would like to pursue if we receive continued funding or the SAFEMITS project.

2.1 Mobile Nodes Abroad

It is important to note that Mobile-IP may be viewed as either an Interior Gateway Protocol or Exterior Gateway Protocol (or neither), simply based on security policies. Put another

way, it is a reasonable security policy to claim that one is not going to allow foreign Mobile visitors using Mobile-IP (or simpler DHCP loaned addresses) to appear "inside" a secure enclave. We previously noted that anti-spoofing measures currently used in the Internet make cross-domain Mobile-IP problematic. We have chosen to implement DHCP-based IPSEC tunnels to demonstrate that Mobile Nodes abroad could securely tunnel home and not have any problems with Mobile-IP source address spoofing (the Mobile Node source address is inside the external IPIP encapsulation and will not be seen until the packet arrives home). However we did not implement any mechanisms that would allow an agent to automatically setup tunnels to tunnel non-trusted Mobile Node packets outside the realm of the secure enclave. This is one mechanism for making smarter security agents that could make cross-domain mobility more feasible.

With such work, the implementers should pay attention to both verification of the security mechanism and should consider risk assessment for what might happen if such mechanisms are breached. One of the problems with dynamically poking holes in "firewalls" is that one may get unintended holes. One of the problems with the notion of "active networks" is that they may be more active than anticipated.

2.2 Intelligent Foreign Agents

As another possible mechanism for dealing with both trusted and untrusted Mobile Nodes, Foreign Agents could be made smarter. We suggest two possible options:

2.2.1 Ad hoc protocols serve as a very primitive but effective mechanism for disallowing cross enclave traffic by Mobile Nodes lacking beacon keys simply because the Foreign Agent would not install a local route for unwanted Mobile Nodes. The end result is that the Mobile Node could send packets but could not receive them through the Foreign Agent. One could improve this mechanism by tying a packet filter access control list to the registration process (802.11 can do this, but at the MAC level).

2.2.2 A Mobile Node could present a certificate to a Foreign Agent, and upon verification, the Foreign Agent would then install an ACL that would permit two-way traffic for the Mobile Node. Note that this mechanism should not stand alone (as it is still "spoofable") from IPSEC measures. For example, a Foreign Agent might also then insist that all forwarded packets must first be send directly to the FA itself using a MN/FA IPSEC association. Packets lacking that relationship would be discarded or tunneled outside the secure enclave. Flexible policy for Foreign Agents may be useful, but appears complex. Of course, the risks inherent in such systems should be considered as well.

2.3 The hard work - integration

It would suffice to say, security software is complex and integration although unloved is extremely important. Any key management system seems to have inherent and possibly untoward complexity that informally appears non-linear in nature. This should be recognized and time should be given to security

and redundancy oriented research projects that takes these problems into account. Given that we are attempting to combine Mobile-IP, IPSEC, redundancy in many forms, ad hoc routing, various operating systems (FreeBSD, Linux), it is no wonder that we have had a difficult integration problem. Our entire design could stand thorough review and slow and patient replacement of key sub-systems. We respectfully suggest that programs oriented towards security and redundancy need to be carefully nurtured and managed in terms of time and budget.

2.4 Keys as a basis for networking

In hindsight our approach to accomplishing an integrated system design is based on a rather curious notion in that we suggested that IP subnets were not the basis of networks. Instead shared trust (keys ...) should be the basis of networking. Our ad hoc systems could be viewed as small groups of Mobile Nodes that formed a network based on individual atomic key knowledge of members. Routing in truth may present a "chicken and egg problem" as you cannot transmit ordinary data before you setup routing as control. But routing protocol security is usually solved by an a priori assumption that a set of routers share a shared secret. We suggest it is not unreasonable to assume that all packets might have a trust relationship and that networking might be done from the ground up (even ARP ...) on that basis.

2.5 Wireless Loading

In the course of the project we discovered that in general non-proprietary wireless drivers do not "load" very well. What we discovered is that the number of simultaneous transfers between multiple laptops and a single agent is strictly limited and certainly does not scale to the level of ethernet. For example, we suspect if all seven members of the project team (assuming all seven had wireless equipped laptops) attempted to simultaneously use FTP to download a large file, perhaps as few as four would succeed in simultaneous download. This is based on an informal experiment but we believe it is truly illustrative of the problem. This has never really been a problem for us due to limited numbers of users. However we expect that anyone who widely adopts a wireless LAN link and then faces > 5 users in the same cell will experience difficulties when simultaneous bulk transfers are attempted. We expect this problem is due to a number of factors including the limited spreading inherent in current spread spectrum techniques (currently limited by FCC regulations), the overall bandwidth, and the fact that such a technology is only capable of "listen while send"; i.e., there is no out of band channel mechanism for true collision detection.

Solutions might lie with techniques borrowed from DARPA "Software Defined Radio" techniques; i.e., a very wide spectrum spreading in terms of frequency. More narrowly based devices might use techniques for sampling limited ranges and then switch to another range that does not show so much current use. An agent could easily include the number of current users in its beacon. More research in this area is necessary.

APPENDIX D

SUNY Institute of Technology at Utica/Rome

Secure Architecture for Extensible Mobile Internet Transport Systems

Eighth Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 December 2001

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Project Status Overview

We began this month operating under the assumption that there will be additional funds forthcoming for FY 02. Under this assumption, we would have enough funding to continue support of our existing students through their scheduled graduation dates. It will also allow modest improvements in our equipment base and support for the remaining graduate students for the next 18 months.

As of now have four M.S. telecom graduate students working on certain aspects of the mobile project at this point. It is conceivable that one of them will graduate by the end of Phase 1 of the SAFEMITS project (May, 2002) and therefore for continuity we have selected others who will remain at least until the end of next year. M.S. telecom student Amitabh Pandey will be graduating at the end of the Fall semester. It should again be pointed out again that should we receive continued funding for another year we plan on making a combined release of Mobile-IP/IPSEC.

Linux Port

Two of our graduate students Rishi Mehta and Shweta Agnihotri are working on porting the mobile-node specific routing daemon code (not agents, just Mobile Nodes) to the Linux operating system. We are finding that many faculty members and students have Linux on laptops and they may be able to make use of this code to implement a Mobile-IP based wireless network.

The remainder of this report will be devoted to those mobile security considerations which are an integral part of the SAFEMITS project.

Mobile Security Considerations

With the rapid growth of Virtual Private Networks, tunneling protocols are assuming a high profile in the Internet. Our work with tunnels as applied to Mobile-IP has uncovered a vulnerability that most tunnels leave unprotected. Basically, while most of today's firewalls stop IP Spoofing attacks, tunnels "drilled" through those firewalls re-enable that class of attack. Strong authentication of the remote systems by the tunnel endpoint, while necessary, is not sufficient to maintain the protection provided by the firewall complex. More generally, if a tunnel server allows authenticated remote systems to become part of a "secure enclave", it must also provide the basic protection that the firewall provides for native hosts in that enclave.

The problem becomes even more interesting if the secure enclave wishes to host "visiting" systems locally. For example, a company might wish to provide Internet connectivity in conference rooms and allow visitors to access the Internet (and not the secure enclave). We will consider these problems below in more detail.

1. Assumptions

Networking, especially when done securely, has been developed from many different perspectives. Each community starts from a presumed base of common language and "normal" assumptions. To minimize confusion, we begin by stating our assumptions and provide a brief description of how commonly used terms are used in this document. We do not mean to imply anything about how they "should be used", just how we chose to use them here.

The focus of this document is on how a secure enclave (firewall protected area) may tolerate Mobile-IP [RFC-2002] or simpler mobility systems (for example, DHCP used standalone) and remain secure. By "secure enclave" we mean a conventional IP site with one management domain and a centralized security administration typically behind one IP firewall [Chapman]. By "firewall" we refer to one or more systems acting together to provide protection for a network. In particular, we assume that one (or more) endpoints of IP tunnels are part of the firewall complex.

Our focus here is on how a secure enclave can protect itself from foreign (non-local) Mobile Nodes. We also deal with IP spoofing issues and possible security problems that might occur due to naive implementations of IP tunneling [RFC-2003] when combined with such spoofing. The discussion is focused on the network layer. We are not considering higher-level authentication or confidentiality services that might be part of an application-level system. When we discuss firewalls, we are mostly talking about network layer access, and such mechanisms as packet-level firewalls with access control, Virtual Private Networks implemented as IP tunnels, and IP layer security (IPSEC) [RFC-1825]. We do not mean to

discourage application or transport layer security in any way rather it is simply not our focus in our research.

Regarding firewalls, we assume Cisco access lists as a rough lingua franca for access control on routers and will use access list examples suitable for Cisco routers. Please see [Ballew] for discussion of Cisco access list mechanisms. We assume packet filter technology simply because accidental holes may indeed be poked through such a router if its manager is not careful.

When we cite IP addresses as examples, we will use private IP addresses as mentioned in [RFC 1918]. These should be viewed as surrogates for public ("real") IP addresses associated with an interior routing domain. We use these addresses because we do not want to cite "real" addresses in any examples.

2. The Problem Space

Firewalls are designed to separate "inside" from "outside". A naive approach to protection would use the source IP address to make the distinction. Unfortunately, IP header information is unreliable as it can be set by the source (or any intermediary) to any arbitrary value. The attacker community knows this well and it forms the foundation for an entire class of attacks known as "Spoofing".

Spoofing has been used as the basis for a whole set of recent attacks (for example, see [RFC-2267]). Some of the attacks are denial of service oriented. Some seek to cause the attacked system to crash or hang. Many of the attacks can be characterized as single packets wherein the IP source and destination addresses in the IP header appear to originate within the attacked systems site.

2.1 Spoofing Attack Examples

To highlight the importance of spoofing attacks, we will briefly discuss three such attacks, TCP SYN [CA-96.21], "smurf" [CA-97.28], and "land" [CA-97.28].

In TCP SYN attacks, the attacker sends TCP initialization packets to a given site. The attacked system is tied up simply due to opening too many TCP control blocks which cause allocation of precious kernel memory. The attacker need only send one TCP SYN packet. The attacker may choose to use a spoofed IP source address so that tracing the attack back to its originating system is difficult.

In "smurf" attacks, the attacker sends one or many ping packets to an IP directed-broadcast address with an IP source address that may also be at the destination site. For example, if a site had a site specific class C address along the lines of 192.168.1.0, the attacking IP destination might be 192.168.1.255 or 192.168.1.0 (0 broadcast addresses may be used as well) with an IP source of

192.168.1.1.

The result is that two systems (at least) may be attacked. The IP source itself is bombarded with ping reflections from all the systems at the directed broadcast address. Further the smurf vehicle could also be used for single packet "ping of death" attacks.

"Land" attacks involve one TCP SYN packet in which the IP source is set to be the same as the IP destination. The attack may cause the receiving machine to hang. In general, note these attacks do not involve packets being returned, unless the packets are returned to another system that is being indirectly attacked itself ("smurf").

2.2 Prevention of Spoofing Attacks

One general technique that can be used is to disallow IP spoofing for internal source addresses [RFC-2267]. Filters can be put in place so that packets arriving on an "outside" interface must have an "outside" source address (or must NOT have an "inside" source address). One may also filter out spoofing attacks attempting to leave from the "inside" of a network.

We will briefly look at how spoofing may be prevented with a Cisco router which we assume is the interface between a site having 172.16.*.* as its internal IP address space and the Internet at large. Note that the access list entries shown here may be part of a more complex firewall policy and/or access list, but we only show the part relevant to IP spoofing.

The following access list entry may be bound to an external router interface and would be applied to packets entering the site.

```
access-list 101 ...  
access-list 101 172.16.0.0 0.0.255.255
```

Note: any packets entering the site with 172.16.*.* addresses will be discarded.

We apply the following access list to packets headed out on the external interface:

```
access-list 111 ...  
access-list permit ip 172.16.0.0 0.0.255.255. any  
access-list deny ip any any log
```

This blocks packets headed out that do not have IP src addresses in the 172.16.*.* range and logs any internal attempts at spoofing. Thus spoofing attacks cannot originate at this site. The result is that a site firewall or border router will neither send or receive packets over an interface when the packets do not belong to the source IP routing domain.

2.3 Anti-Spoofing Measures and Mobile-IP

The result of such anti-spoofing measures is that packets headed into the enclave to "foreign" Mobile Nodes; i.e., Mobile-Nodes from some other site than 172.16.*.* will not be permitted to enter. Packets trying to leave the site from systems from another site, say 192.168.1.0, will not be permitted to leave. Mobile-IP within the same routing domain, which we might call "interior Mobile-IP" would be permitted. Mobile-IP ("exterior Mobile-IP") between two interior routing domains would not be permitted.

Now we must consider what happens to packets sent from a "visiting" foreign Mobile Node that is somehow operating within the secure enclave. First the UDP-based Mobile-IP registration protocol would still work if Foreign Agents were used as Foreign Agents act as UDP proxies for Mobile-IP registration; i.e., they will replace a Mobile Nodes IP source address with their own (legal) source address. A Mobile Node using DHCP as a source of local addresses could also succeed if it used the DHCP-obtained local address.

Data packets sent directly out of the domain from the visiting Mobile Node (unless tunneled via a local IP source address) would be discarded at the border. Data packets that somehow escaped the local secure enclave's border router could also be discarded by the "home" border router's spoofing filter as well, as it would not permit packets to enter that have "local" IP source addresses.

Data packets tunneled from the (exterior) Home Agent to the (interior) Foreign Agent would be allowed through because the external encapsulation would get past the spoofing filter; i.e., Home Agent to Foreign Agent IPIP packets would have the legal interior IP source for the Foreign Agent as the exterior IP source address.

Home Agent <-----> Border Router (for 192.168.1.X domain)
(192.168.1.X)
^

| the Internet

|

v

Border Router (secure enclave)

^

|

| secure enclave (172.16.*.*)

v

2221

Foreign Agent (172.16.*.* address for Care Of Address)

^

|

v

Visiting Mobile Node (192.168.1.X)

In addition to the obvious problems that anti-spoofing raises for Mobile-IP, one must also ask if tunnels raise additional security concerns and how one might address both those concerns and security for both the Mobile Node itself, its home domain, and the "visited" domain too.

3. Tunnels Considered Harmful

One mechanism that is part of Mobile-IP and in point of fact many other routing protocols are IP tunnels which might be implemented with IPIP, IPSEC Tunnel Mode, or Cisco's Generic Routing Encapsulation [RFC-1701]. It should be pointed out that IPIP tunnels are not peculiar to Mobile-IP. They are used in many routing protocols for many purposes including tunneling non-IETF protocols or building virtual networks on top of the current Internet.

Tunnels may mean encapsulated packets where one has one IP datagram inside another IP datagram and we will use IPIP (IP protocol 4) as our example here. Mobile-IP uses tunnel mechanisms like IPIP to forward packets from the Home Agent to a remote "Care Of Address". The COA is a local site IP address that may represent a Mobile Node that has acquired a local IP address itself directly via DHCP or a router system that understands Mobile-IP called a Foreign Agent. Any Mobile-IP system, including Mobile Nodes, Home Agents, or Foreign Agents, may source or sink tunnel packets.

When a Home Agent forwards packets to a Mobile Node that is at a Foreign Agent, the use of IPIP in a datagram may appear as follows:

IP outer header IP inner IP datagram

ip src= Home Agent ip src = peer end host | TCP, etc.

ip dst= Foreign Agent ip dst = Mobile Node |

|

| packets to MN

v

Home Agent ===== IPIP tunnel to COA ==> FA and Mobile Node

One might ask if it is enough to simply use IPIP tunneling and somehow tunnel either from the Foreign Agent or Mobile Node back to the Home Agent and thus evade the anti-spoofing measures at a firewall? Unfortunately, this is an insecure approach. In the first place, it is not enough to simply tunnel over the IP spoofing firewall. This is simply a new form of spoofing which we might call: "IPIP spoofing". The problem is that if one has a tunnel sink (be it any kind of agent or Mobile Node) that "decapsulates" packets and then forwards them, others can launch their spoofing attacks with IPIP too and thus have the spoofing emerge "inside" the enclave firewall. For example, smurfing might simply be

redirected through a tunnel. The inner IP header might be directed broadcast with an interior IP source named as a target. All the previous attacks (TCP SYN, "smurf", "land") can thus be done through the firewall.

We suggest that one can block tunnels with current access list mechanisms and thus control tunnels so that tunneling from the outside can only be done to certain hosts that will be considered as "network-layer" bastion hosts.

For example, with Cisco IOS one can add the following statements to access list 101:

```
access-list 101 deny ipinip any any
access-list 101 deny gre any any
```

thus blocking any IPIP or GRE packets coming in over the router. One may further add a permit statement to force any IPIP packets coming in to land at a certain host and then treat that host as a bastion host; i.e., a nexus of security focus.

For example, in a Cisco input access list, means IPIP will only be allowed to the host 172.16.1.3, which we will assume is a tunnel sink agent.

This can be accomplished using the following commands:

```
...
access-list 101 permit ipinip any host 172.16.1.3
access-list 101 deny ipinip any any
access-list 101 deny gre any any
...
```

We have focused internal trust for IPIP on that one system. We next need to explore how to make sure that packets arriving at the tunnel sink agent are **not** attacks that can be made via a tunnel sink. This can be done with "tunnel-mode" IPSEC tied to IPIP. We will discuss this idea further in the next section.

4. Problem Solution Space

We will discuss how to solve these problems from two topological points of view. First we look at the situation from the Mobile Node abroad's point of view. We assume it actually wants to get packets home and not compromise home security. Thus this point of view must necessarily include the Mobile Node's home enclave. We then look at the situation from the "foreign" security enclave's point of view. We assume that the foreign enclave wants to allow mobile service but protect itself. We also must consider the question of how both security enclaves (home and away) in general protect themselves from any tunnel-based attacks.

In this discussion we also try to contrast the use of Mobile-IP versus a simpler form of DHCP-only mobility that does not use Mobile-IP. Keep in mind that the main semantic for the use of Mobile-IP is that the Mobile Node retains at least one fixed IP address that is non-local for the subnet it is visiting. A Mobile-IP system may have two IP addresses associated with it, a fixed permanent Mobile-IP address (call it the "Mobile-IP address"), and a locally acquired address (call it the "DHCP address"). A DHCP-only system would only have one locally acquired IP address.

4.1 Mobile Nodes Abroad Point of View

In this section we will consider the problems for a secure enclave if Mobile Nodes in that secure enclave go abroad; i.e., out to the Internet beyond the firewall. We must ask how the Mobile Node can secure its own traffic and in effect, take its security enclave with it. We must also ask how the home enclave can secure traffic coming from that Mobile Node back inside, which will extend the thinking about tunnels in the previous section.

Glass and Gupta suggested that Mobile-IP Mobile Nodes abroad may use DHCP to acquire "local" IP addresses, Thus they can get by the anti-spoofing measures in the firewall router. This is indeed a reasonable possibility. Further, the Mobile Nodes can use IPSEC with two-way tunnels between the Home Agent as a classic bastion host and the Mobile Node.

4.1.1 Packets from the Mobile Node Out

If we assume Mobile-IP and two addresses in use by the Mobile Node, packets tunneled from the Mobile Node to the Home Agent might have the structure:

IP(1) | IP(2) | <IPSEC> | IP(3)

Each IP header has its own purpose. The most external header, IP(1) exists to get the packets to the Home Agent with the DHCP acquired address == 172.16.1.2. The IP destination would be 192.168.1.1. Thus header(1) allows transit of the Internet and any anti-spoofing firewalls. When the packet arrives at the Home Agent, that header is discarded and header(2), consisting of IP(2) combined with an IPSEC header is processed. Here we assume that the Mobile-IP address 192.168.1.2 is used for the source address and the Home Agent is again the destination.

The fixed Mobile-IP address may be needed here as it allows a-priori manual IPSEC keys to exist between the Mobile Node and the Home Agent. In effect, this is an IPSEC tunnel between the Mobile Node and the Home Agent. The interior header would contain the Mobile Nodes fixed address (192.168.1.2) as IP source and the address of any destination to which it is allowed to send packets.

The aforementioned triple-header system could be optimized by a higher-level

protocol that could produce a dynamic binding between the local DHCP-acquired COA and the Home Agent's destination address. There is no reason Internet Key Exchange protocols [IKE] where non-IP naming schemes are used could not be deployed here. This would allow one header to be deleted.

For a DHCP-only form of mobility, the packet layout situation would be simpler. The Mobile Node would use a non-IP naming scheme with IKE to form a security association with a Home Security Agent. IP header (2) would not be needed.

4.1.2 Packets Coming From Home to the Mobile Node

For Mobile-IP, we must now consider packets coming back to the Mobile Node via a tunnel from the Home Agent. By definition these packets are tunneled to a local IP address and are not subject to problems caused by anti-spoofing filtering. However IPIP unadorned is a security threat to the receiving enclave. And of course, the Mobile Node may choose to have IPSEC-based security between itself and its home enclave.

One possible encapsulation scheme might take this form:

IP(1) | IP(2) | IPSEC | IP datagram

IP(1) exists to get the packets from the Home Agent to the remote Care Of Address which might be a Foreign Agent or a Mobile Node that has acquired a local IP address. The inner IP header would exist where manual keying is needed with IPSEC and the IP source would be the Home Agent. The IP destination would be the Mobile Node itself. Note that again IKE could be used to optimize out an IP header as long as IP addresses are not part of a manual configuration scheme.

It is highly likely that from a security policy point of view, one would not form security associations (especially confidentiality-based security associations) between random Home Agents and random enterprise-external Foreign Agents. As a policy consideration, unsecured IPIP might simply not be allowed to Foreign Agents.

Foreign Agents might insist that all IPIP packets be sent to them from internal Home Agents with which they share an a priori security association. Alternatively Foreign Agents might exist "outside" a secure enclave, or unadorned IPIP packets when decapsulated might only be allowed to go "outside".

4.1.3 Tunnel Security at Tunnel-Exit Agents

We suggest that a tunnel-sink agent like a Mobile-IP Home Agent may want to guarantee that all packets sent to it via a tunnel are cryptographically verified; e.g., shared secret keys might exist between it and the Mobile Node abroad. No packets forwarded to the tunnel-sink agent by the firewall will be internally decapsulated and forwarded until they have been cryptographically verified.

This might be accomplished with an access list mechanism tied to IPSEC or by simpler means. For example, the PSU system mentioned above has a BSD sysctl(8) switch:

```
# sysctl -w net.inet.ip.mvifipsecinput=1
```

That if set forces the IPIP driver to only forward packets if and only if IPSEC authentication or decryption has successfully occurred between the remote system and this system. As a consequence, one may be sure that a Mobile-IP Home or Foreign Agent or any tunnel sink only forwards IPIP packets that have successfully passed IPSEC processing. Put another way, a security association must exist between the tunnel sink and the tunnel source system.

Packets coming from remote security-aware Mobile Nodes might have several forms:

IP(1) | IPSEC | IP datagram

or possibly

IP(1) | IP(2) | IPSEC | IP datagram

For example, the former packet architecture might occur with a remote Mobile Node that is only using DHCP and wants to securely tunnel home. The latter might be used by a remote Mobile Node that is using Mobile-IP and has also used DHCP to acquire a local COA. The local anti-IP-spoofing firewall might then be configured in a number of possible manners depending on local security policies and the structure of external but acceptable packets.

For example, with current Cisco access list technology, we could permit IP | IPSEC packets using ESP (ip proto 50) or AH (51) to the Home Agent:

```
...
access-list 101 permit 50 any host 172.16.1.3
access-list 101 permit 51 any host 172.16.1.3
access-list 101 deny ipinip any any
...
```

As in our previous example, the firewall might simply allow IPIP but only to a Home Agent. This would apply to the second IP | IP | IPSEC example. We must point out that the security problems here are not terribly different from those encountered by current dialup clients into a secure enclave that access the enclave via an internal terminal multiplexor. The exterior host tunnels into a secure enclave and an agent in the secure enclave applies cryptographic measures to packets that have come in from the outside.

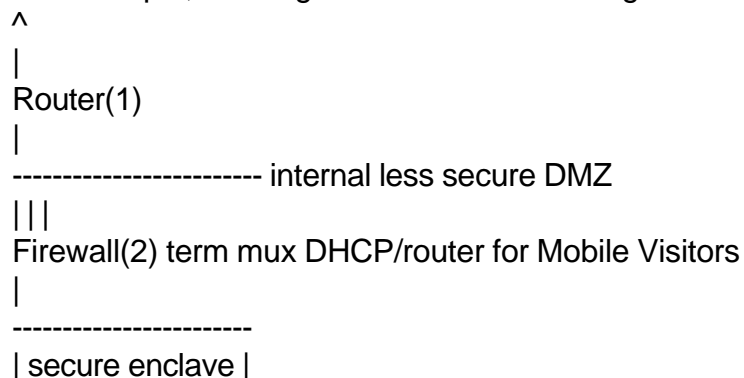
4.2 Hosting Visitors From Abroad

In the previous section we have focused on how to protect the Mobile Node abroad and also discussed the problem of how to make tunnel (exits) more secure. One must also worry about the security of the "other" enclave, else enclaves may not desire to host foreign Mobile Nodes. It makes little sense for a firewall-protected enclave to allow visitors to penetrate the enclave at will and thus enable possible attacks on internal systems by visitors.

Of course, we could start with a security policy that does not allow visitors to penetrate the firewall. In effect, that is the current security policy for many sites. However it is our goal here to discuss how we might tolerate "less trusted" visitors, not define them out of existence.

We suggest a topological approach based on network design measures that can be made with current (or near-current) technology and that should allow a secure enclave to remain secure. Our basic principle is: "design the network so that visitor packets are not allowed inside". We observe that whatever is done to implement this goal will probably be similar to current systems that have two-level security enclaves.

For example, one might have a network designed as follows:




The above may be viewed as a logical (not necessarily physical) structure. We have a Router(1) that simply serves to allow access to the Internet. Inside the external router we have a less secure DMZ network that may serve to allow unfiltered access to the Internet. This network might include terminal multiplexors and local mobility servers. Behind it we could have at least one level of firewalls with bastion hosts (which may or may not be on the DMZ network). Firewall(2) would serve the function of the traditional firewall machine.

One might observe that the above scheme does not force visiting laptops to acquire local addresses to bypass IP-spoofing filters. True, but unfortunately there may be other firewalls on the way home that possess such filters. Certainly the firewall at home may possess such a filter. The above may be viewed as a physical network design that would tolerate visitors. However it is very likely that such a design may not be very convenient to implement. We suggest that various virtual network techniques may be used to approximate the physical structure above.

Let us briefly consider two methods that may be used to logically separate networks and thus remove construction difficulties and make the deployment of mobile networks easier. Note that these techniques **require** careful network security design. The security onus here is certainly on both the network designer, and on the creators of both Mobile Agents and security gateways. Caution needs to be employed in the design of such systems and functionality must be clearly communicated by implementers to potential network designers.

First one may use a combination of simple tunneling sans IPSEC and/or authenticated tunnel packets (e.g., IPSEC AH) between a mobile router (MIP foreign agent or DHCP router/server) and the "outside" network. The basic idea would be that a router (e.g., a Mobile-IP Foreign Agent) could take all packets presented to an "exterior" interface and tunnel them (using IPIP or GRE) (possibly with additional IPSEC to alleviate paranoia) to a firewall-exterior interface on a border router. As a result, visitor packets would have no opportunity to access interior hosts. They would be tunneled "outside" and would be treated as external packets coming back through any existing firewall mechanism.



Note that we assume here that the agent(3) and the exit tunnel router(2) are under the control of the same network administration.

We suggest that careful combination of access lists with tunnel technology should allow the above picture to be collapsed in various ways. For example the Firewall(1) and exit router(2) systems could be the same system. In addition, the router(3) that enables mobility could potentially optimize packet delivery. If IPSEC security associations existed at that router between a Mobile Node and the router itself, it might choose to NOT forward IPSEC-verified packets that show up via the external visitor link(4) over the internal tunnel. Thus IPSEC packets from "local" mobile systems that belong to the enclave itself could be allowed direct access to the local enclave. Of course, packets that lacked a security association with the mobile agent router would be forced over the tunnel to the "outside" world.

We will not go into details, but link-layer switching technologies can also be of use here. For example, Virtual Lans [3com] when assumed to be 1-1 with IP subnets could be used as a way to funnel visitor packets back to a router that might apply access list technology to packets trying to cross from an "exterior" subnet to an "interior" subnet.

5.1 Miscellaneous Considerations

Although we cannot attribute such discussion, some have suggested that some sort of Firewall Discovery system might allow Mobile Nodes to dynamically tunnel to and from firewalls. There are several problems with this notion:

1. It is unnecessary since our solution here will work with current or near-current firewall technology.
2. It is not very likely from a security point of view.

Security people and network managers may not care for notions that involve poking holes dynamically through firewalls. Complexities involved in cross-security domain certification are likely beyond near-term scope. Further the security folks "at-home" may not care for schemes that involve key exchange with strangers; i.e., a Mobile Node from home somehow secures packets between itself and a foreign firewall at a different enterprise. After all, that firewall might choose to store all data traffic, and enable a classic "man-in-the-middle" attack.

3. Traditional notions of IP fate sharing (considered unacceptable) may apply here. Mobile-IP systems are already tied to the fate of their Home Agent. Additional ties between systems that are not related from internal routing or security enclave considerations may be complex. After all, it is hard to predict how many firewalls that rule out IP spoofing to/from a given site may exist. Schemes that allow trusted locals to poke holes through firewalls are perilous by definition since "untrusted" people may crack the scheme. It is unlikely that dynamic mechanisms that allow random visitors access will prove widely acceptable.

5.2 The Role of a Mobile-IP Foreign Agent

In the previous discussion, we suggested that DHCP can be used to simply allow Mobile Nodes abroad to obtain a local address. Using that address they can then send packets wherever they choose. As a result, it might seem that there is little role for a Mobile-IP Foreign Agent in a security system. Ultimately the roles that mobility systems play depend upon policy considerations. One could suggest a policy wherein Mobile Nodes abroad are not allowed to talk directly to (as opposed to through) or exchange cryptographic material with "foreign agents". This is certainly a reasonable policy. However the focus of such a policy is on the Mobile Node.

We need to also consider Foreign Agent oriented policy and how a Foreign Agent might serve as a border router for a secure enclave. Foreign Agents may serve as routers that

simply do not allow foreign visitors any access to an internal enclave and only allow authorized local Mobile Nodes entrance. Many techniques exist for such screening including the pre-existing Mobile-IP manually keyed registration that can secure Mobile Node access via a given Foreign Agent. However, security techniques should apply to all packets and not just Mobile-IP registration packets.

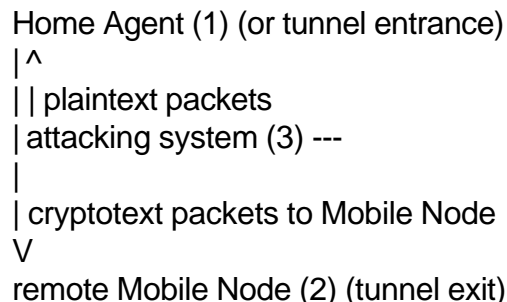
As another possibility (and there are probably others), "IPSEC-aware" Foreign Agents can discriminate between locals (who hold security credentials) and non-locals (who lack security credentials). Once a Mobile Node has identified itself to a Foreign Agent as belonging to the agent's secure enclave, it could use an IPSEC security tunnel between itself and the Foreign Agent. Any packets verified by the Foreign Agent as belonging to the local secure enclave could thus be delivered locally and not pushed out of the routing domain via a tunnel. Non-local visitor packets might be unceremoniously escorted off the premises via another kind of tunnel and would have no access to internal resources. Thus local mobile systems and visitors could both be tolerated at the same agent link.

Of course a paranoid enclave might choose for policy reasons to force all wireless visitors to be "foreign". "Locals" could be always treated as remote visitors and tunneled outside, thus having to use secure means to come back inside. Or foreign agents might simply not be permitted entrance at a given agent. Both policy considerations are possible and should be considered in implementations.

6. Security Considerations

Our research focuses specifically on issues arising out of the interaction between firewalls and any tunneled protocol and highlight security concerns regarding Mobile-IP or simple DHCP for foreign visitors "beyond" the home firewall. We should point out at least one more specific security consideration for tunnel entrances. If IPSEC is used in "tunnel-mode" at a router or forwarding system that is neither the IP source or IP destination, it is possible that the security system may be subject to "proposed plaintext attacks".

Refer to the following figure:



If an attacking system(3) can present plaintext packets to (1), and then read them back after encryption in the tunnel between (1) and (2), the potential for proposed plaintext attacks exists. This liability exists for a number of proposed combined tunnel and security systems,

as long as network-layer forwarding combined with IPSEC (or cryptography) is part of the architecture. Solutions for the problem include session keys [IKE] and possible restriction of communication between the Home Agent(1) and the Mobile Node(2) to exclude IP sources that do not lie within the home enclave. By definition, this problem is found only with network-layer forwarding (i.e., at IPSEC gateways), and is not present in any end-system to end-system IPSEC.

7. Conclusions

In this status report we have presented proposals that will enable Mobile Nodes from abroad or nearby to less insecurely access the Internet. Such systems are not dissimilar from current dialup systems that involve a remote PPP-based dialup client and a local terminal multiplexor. IPSEC-enabled tunnel mechanisms may be used between the Mobile Node system and its home security companion. Very simply put, the Mobile Node is an extension of the local security domain. However, in addition to securing the Mobile Node and its home enclave, one must also give thought both to the dangers of tunnels and to how a local enclave may enable its own security and still tolerate visitors.

In summary, we will make the following suggestions:

1. DHCP to acquire a local COA solves problems caused by IP spoofing prevention for visiting Mobile Nodes abroad and may or may not be combined with Mobile-IP.
2. Suitable two-way cryptographic tunnels between a system abroad and a routing system at home will allow a Mobile Node's own traffic to be securely tunneled over the Internet.
3. IPIP tunnels sans cryptographic safeguards should be viewed with caution. If an IPIP tunnel sink does not guarantee cryptographically controlled access, an attacker may tunnel various one-way attacks (land, etc.) into an enclave. The tunnel sink may be logically regarded as an extension of the firewall itself. It may be co-located. If it is not co-located, firewall filtering mechanisms may need to be duplicated at the tunnel-exit point.
4. Flexibility in routing, access list mechanisms, and encapsulation possibly with authentication should be considered by implementers so that a secure enclave can securely escort visitor packets off-site without threat to secured systems within the site.
5. Security considerations must apply both to Mobile Nodes abroad, their own home enclave itself, and also to how enclaves may be designed to tolerate visitors.

References

Atkinson, R., "Security Architecture for the Internet Protocol", August 1995.

Ballew, Scott, "Managing IP Networks", O'Reilly and Associates, Inc., 1997; ISBN 1-56592-320-0

Bellovin, Steve, "Report of the IAB Security Architecture Workshop", April 1998.

Chapman, D.B., and Zwicky, E.D., "Building Internet Firewalls", O'Reilly and Associates, Inc., 1995

Ferguson, P., and Senie, D., "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", January 1998.

Gupta, V., Glass, S., "Firewall Traversal for Mobile IP: Guidelines for Firewalls and Mobile Ip entities", draft-ietf-mobileip-firewall-trav-00.txt, work in progress,

Hanks, S., Li, T., Farinacci, D., Traina, P., "Generic Routing Encapsulation", October 1994.

Passmore, David, and Freeman, John. "The Virtual Lan Technology".
3com Inc.

Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets",

[CA-96.21] CERT Advisory CA-96.21; TCP SYN Flooding and IP Spoofing Attacks; September 24, 1996.
http://www.cert.org/advisories/CA-96.21.tcp_syn_flooding.html

[CA-97.28] CERT Advisory CA-97.28; "Teardrop/Land" IP Denial-of-Service Attacks; December 16, 1997.
http://www.cert.org/advisories/CA-97.28.Teardrop_Land.html

[CA-98.01] CERT Advisory CA-98.01; "smurf" IP Denial-of-Service Attacks; January 5, 1998. <http://www.cert.org/advisories/CA-98.01.smurf.html>
<http://www.3com.com/nsc/200374.html>;

APPENDIX E

SUNY Institute of Technology at Utica/Rome

Secure Architecture for Extensible Mobile Internet Transport Systems

Seventh Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 November 2001

Accomplishments

In this report we present our plans for deploying a combined layer 3 Mobile-IP and IPSEC routing architecture. We discuss possible routing security architectures and then present two alternative designs for an integrated Mobile-IP/IPSEC routing architectures. We refer to these as "closely-coupled" and "loosely-coupled". The closely-coupled architecture relies on a direct binding of IPSEC policy attributes to routing table entries by Mobile-IP routing daemons. The loosely-coupled architecture is based on a more traditional access control list association between the FreeBSD 4.3 Mobile-IP/IPSEC Implementation and a Linux based Mobile-IP/IPSEC integration. Our discussion concludes with an architectural analysis of combined Mobile-IP/IPSEC and a call for the use of IPSEC as part of any mobile VPN scheme.

Introduction

Recently wireless security as found in the popular 802.11 [10] wireless protocol has suffered a series of failures due to the apparent collapse of part of the 802.11 specification called WEP, for "Wired Equivalent Privacy". WEP was intended to offer security services including encryption and authentication. Security experts have recently demonstrated substantial security problems with WEP. Borisov, et. al.[4], describe security and network architecture flaws in WEP. A recent cryptanalytical paper [6] then described a theoretical attack against the RC4 stream cipher used in WEP. Worse, in a recent ATT technical report [17], the theoretical attack was implemented.

The ATT paper in its conclusions suggests that the 802.11 link layer be viewed as insecure. To be fair, the IEEE 802.11 specification stated that WEP was not only optional, but was intended to make wireless "at least as secure as a wire". Unfortunately, this may be misleading to naive users who assume that WEP would offer serious confidentiality services. Borisov's paper, and the ATT technical report both suggest that users should consider using higher-level protocols; for example, IPSEC [9] and Secure Shell [18]. We concur and further suggest that IPSEC could be directly combined with Mobile-IP [12] in order to make a Virtual Private Network mechanism that is completely based on the layer 3 network layer. This idea was initially proposed by DARPA funded research along those lines at Portland State, between 1995-1999 [15]. In the SAFEMITS project, we intend to explore two different experimental system architectures for a combined Mobile-IP/IPSEC implementation. It should be noted that the earliest architecture was created on the FreeBSD platform using an IPSEC, originally done for NetBSD by the Naval Research Labs, and the PSU Mobile-IP implementation. Our architectural design will be based on the current Linux operating system. We intend to use an existing port of Mobile-IP which was designed to work with Linux to make an integrated IP SEC/Mobile-IP.

Mobile Routing Security Policies

During the first phase of the SAFEMITS project we decided that from a formal point of view, we could distinguish three very general architectural frameworks for mobile routing security. We will call these architectural constructs secure mobile routing architectures. These architectures are as follows:

- MN to security gateway VPN: A two-way VPN is setup between a Mobile Node and some security gateway that acts as an "entry point" into a secure enclave. From the point of view of traditional firewall thinking, the security gateway is a bastion host. In Mobile-IP terms, it may be co-located with the Home Agent (which is the assumption we make in our implementation). Using IPSEC, we setup a two-way layer 3 ESP tunnel, which might or might use dynamic keying. As we will present later in more detail, we have implemented this form of VPN in both of the aforementioned Mobile-IP/IPSEC implementations. Of course, other possible VPN technologies may be used. A Mobile Node outside a secure enclave, has a two-way IPSEC VPN to and from its Home Agent. Foreign Agents are not involved directly in any security association and are merely tunneled over (as are any other layer 3 entities). The first link may be assumed to be wireless, and can be assumed to be outside the secure enclave. The path between the MN and security gateway may be multi-hop and may span the Internet or barring the first link, may be internal to the secure enclave.

In terms of the number and scalability of key associations, key management is linear; that is, for each HA, we have a set of MNs. Key management may be made more complex by security gateway (HA) redundancy issues. We do not rule out a centralized key management system within the secure enclave, that might, for example, use DNS or some other system.

-Agent boundary VPNs: In this form, we restrict cryptographic services to the "external" link; that is, MNs are assumed to be outside the secure enclave, have two-way VPNs between themselves and a boundary agent, and the link connectivity is confined to only one link. The typical boundary agent could have one external link and one internal link. Boundary agents might be layer 3 entities as with Mobile-IP agents, or layer 2 entities as with 802.11/WEP access points. Boundary agents in a MIP system would insist that MNs must have an a-priori security association. Thus MNs that do not have local IPSEC keys would not be able to penetrate the secure enclave security architecture, FAs, by definition must belong to your security enclave, and MN-FA security associations must exist. Note that in the previous architecture, FAs were not part of the picture. Manual key administration here is fundamentally not scalable. as key associations are a function of the number of boundary agents times the number of Mobile Nodes. We believe that an internal tie-in between IKE daemons and centralized key service, possibly via DNSSEC is mandatory. For example, the Portland State University project did not implement such an architecture, although they did view it as possible future work. However, the PSU researchers did implement a layer 3

authentication system for Mobile-IP itself, that required authenticated ICMP advertisements from all network elements including agents and MNs [2]. Both BBN [19] and SMN also implemented layer 3 authentication systems based on per node digital signatures. Note that such authentication systems are not intended as replacements for higher-order confidentiality systems like IPSEC. They are merely supplemental.

- Secure multi-hop ad hoc routing: Multi-hop ad hoc routing refers to Mobile Nodes that setup multi-hop routing paths via a new class of dynamic routing algorithms; for example, please see [3]. The PSU project implemented a form of DSR [7] in which end host to end host IPSEC associations were manually available. Thus all packets between any two MNs could have IPSEC applied to them. It is important to note that "consenting" MNs in such an architecture, by definition, belong to the same security domain.

In the following section, we present our design of a closely-coupled IPSEC/Mobile-IP architecture. Next we present the alternative loosely-coupled architecture in which we have combined our Mobile-IP with IPSEC. In section 5, we present some architectural analysis in terms of system organization, and finally we present our conclusions.

Closely-Coupled Linux Mobile-IP/IPSEC

We will briefly discuss some architectural aspects of our combined Mobile-IP/IPSEC architecture. Much of the Mobile-IP architecture itself has proved portable over time, but the IPSEC aspects themselves did not survive abandonment of the Naval Research Labs (NRL) IPSEC mechanisms, based on older RFC 1825 IPSEC.

We will briefly describe a new IPSEC mechanism based on close coupling of routes and IPSEC policy. We call this closely-coupled because the route daemons directly manipulate the IPSEC policy. Our design calls for creating an experimental IPSEC policy system based on modification of the route(4) socket.

IPSEC assumes two abstract databases in the operating system, that can be used for cryptographic operations on packets. One may be called a policy database with rules similar to: use IPSEC (tunnel/transport/ESP/AH), on these IP addresses, with a certain security association (algorithms/keys). Formally this database is called the Security Policy Database or SPD. The other database provides key material; for example, use BLOWFISH, 3DES, with certain key bit strings, and is called the Security Association Database or SAD. Our design for routing socket modifications will allow routes in the routing table to act as the SPD. We assumed key material had a priori been loaded into the SAD. Thus the SPD references the SAD for actual key materials.

Logically the route(8) command could be assumed to have the following form:

```
# route <ipsec-mode> -spi <SPI> -itsrc <SA-ipaddr> -itdst <SA-ipaddr>
```

The ipsec-mode could be any of -ah, -esp, -ahtunnel, -esptunnel. The modes defined a particular route as either transport or tunnel mode IPSEC.

When a route was loaded, either manually or by a mobile routing daemon, internally a search was performed in the kernel for the SAD, and if an appropriate binding was found, a pointer was setup between the routing table, and the SAD.

The Linux operating system has chosen to adopt the version of IPSEC developed in Helsinki, Finland. This system is based on the RFC 2053 version of IPSEC. Of course, it has also never been burdened with US export law problems.

It should understand that most conventional IPSEC implementations are based on rulesets similar to firewall access control lists. The Mobile-IP/IPSEC routing table feature was used for a number of different security architectures. For example, we plan on implementing the basic MN to security gateway VPN. The HA to MN routing path will require creating an ESP tunnel on the IPIP tunnel device, resulting in packets with an IP ESP (IP datagram) header structure. Other security features will include HA to FA 1-way authenticated tunnels with a IP AH IP datagram structure as opposed to the conventional IPIP tunnel. Also our mobile-node daemon will be capable of using a combined form of DHCP, ESP, and Mobile-IP, when no foreign agents are to be found. This design allows a mobile node to retain its invariant MN IP when away from its home IP address area. The use of the DHCP IP address as the COA meant that any possible IP ingress address problems were avoided because the COA address did not belong to the MN's home addressing domain (see, for example [1], and [5]). Thus, Mobile-IP enabled systems can wander away from their home security enclaves without having to worry about the IP source ingress filter problem.

Loosely-coupled FreeBSD Mobile-IP/IPSEC

The current architecture, based on 4.3 FreeBSD, combines the KAME IPSEC implementation and the PSU Mobile-IP daemons. We have re-implemented the basic MN to security gateway VPN. We refer to this architecture loosely-coupled because the Mobile-IP daemons do not directly manipulate the IPSEC policy. In KAME, IPSEC policy is setup more on the lines of traditional layer 3 access control lists. We assume initial IPSEC two-way tunnels are setup between the Home Agent and Mobile Node, and then run Mobile-IP on top of that configuration. In this section, we will explain the implementation setup in detail, and discuss some resulting implementation problems and solutions.

From the high level point of view, as routing consists of two 1 way problems, we must deal with 2 problems, 1. MN to HA, and 2. HA to MN. For IP datagrams sent from the MN to the HA, we tunnel conventional IP datagrams from the MN to the HA. Thus the IP outer header has an IP src = MN IP, and an IP dst = HA IP. The ESP header encrypts the interior datagram sent from the MN to some other host. For the HA to MN path, we first have packets tunneled via an IPSEC tunnel (IP ESP, IP datagram), where the outer IP header has an IP src = HA IP, and IP dst = MN IP. This packet is then encapsulated inside an IPIP datagram that deals with the COA. Conceptually the HA to MN path can thus be viewed as (IP (dst=COA), IP ESP, IP datagram).

Architectural Details

The FreeBSD 4.3 KAME/IPSEC system allows three levels of kernel control (see `ipsec(4)`). `Sysctl(8)` can be used for global policy. The manual `setkey(8)` command is used to set IPSEC packet-filter defaults { which are similar to traditional layer 3 access control lists implemented in routers. In addition, `setsockopt(2)` can be used for setting per socket IPSEC policy attributes. Thus routing daemons could choose to avoid more general policy when warranted. We make no use of the `sysctl` mechanism and instead use a combination of the `setsockopt(3)` and `setkey(8)` mechanisms.

Mobile-IP interoperation

The basic MN to HA two-way VPN policy requires several modifications to the Mobile-IP daemon implementation. First of all, as one possible security policy choice, we chose to make the necessary UDP and ICMP sockets, bypass any and all IPSEC packet mechanisms in the kernel. This is done using `ipsec set policy(3)`, and `setsockopt(2)` calls. This means that all Mobile-IP packets bypass local IPSEC, and must rely upon their own devices for security. Remember that we choose to ignore Foreign Agents, thus it is important that we be able to talk to them and not assume we speak IPSEC with them. Further, by definition, we cannot share secrets with agents from another security domain. Hence we choose to let Mobile-IP as a protocol stand on its own, otherwise MNs would wrap Mobile-IP registration packets in ESP, FAs might not understand them, and thus could not relay them to the Home Agent.

The second implementation aspect is unfortunately far trickier. When a MN visits a FA, "all" packets in theory will be delivered via a HA tunnel encapsulation; that is, datagrams processed by the KAME IPSEC tunnel are formed as IP ESP f IP datagram g, with the outer IP src = MN IP, and the outer IP dst = HA IP. Unfortunately this runs full tilt into the BSD ARP table implementation. In the current BSD architecture, the ARP table is not separate from the routing table, and is implemented via a so-called clone route mechanism. When an interface uses ARP, and its IP address is configured, a clone route is placed in the routing table. For example, in the sample routing table below, we see a clone route (marked with the UC flags) that was loaded for a local ethernet interface when the interface was booted. One ARP table entry was instantiated in the routing table for local IP address 10.0.0.1, and later filled in by the ARP protocol itself with the MAC address of the local link host, 10.0.0.1.

```
host# netstat -rn
Destination Gateway Flags Netif Expire
10.0.0.0/8 link#1 UC 0 0
10.0.0.1 0:d0:c0:5b:18:0 UHLW 4 3
```

This means that when a MN visits a Foreign Agent, and the first packet is sent via the IPSEC ESP tunnel from MN to HA, the outer IP header will of course, have IP dst = HA IP. This in turn, will cause the clone route to create an ARP table for the HA, because the MN after all, shares a local IP subnet association with the HA. Naturally since the HA is not nearby, this causes complete failure as no packets can reach the HA.

In order to fix this problem, we will need to modify mnd to take advantage of the state machine. When configured for IPSEC, and in NOWHERE or AWAY states, it simply deletes all ARP table entries, and also deletes the clone route. When at home, the clone route is reinstalled. This is one possible policy choice, and the implementation might eventually allow more flexible configuration policies.

Architectural Analysis

In this section we wish to present an architectural analysis and briefly consider two questions:

1. what key ideas might be necessary in an operating system architecture to allow a combined Mobile-IP and IPSEC? and
2. What are the pros and cons of the two Mobile-IP+IPSEC architectural approaches, themselves?

We suggest that KAME IPSEC has provided us with two necessary features that we hope would be available with any IPSEC implementation. The first feature that was important is the ability to specify with the KAME packet filter mechanism that "all packets" should be sent over a tunnel to a tunnel endpoint. For example, the MN should be able to send "all" packets to the HA. It is hard to imagine that an IPSEC implementation would not have this capability, but it is fundamental and necessary.

The second important IPSEC capability is the ability to override higher-level "all packets must use IPSEC" packet filters on a per-socket basis. Without it, Mobile-IP registration packets could not be relayed by Foreign Agents that do not belong to the security domain. More generally, it is extremely reasonable for routing daemons using any routing protocol to be able to except themselves from a system-wide IPSEC policy. Most routing protocols have their own authentication mechanisms; for example, OSPF [11] has per link authentication.

We have arrived at the conclusion based on our preliminary analysis that the "Helsinki" Linux based IPSEC/MIP experimental implementation was strongly-coupled, because the IPSEC policy was directly manipulated by the mobile routing daemons. On the other hand, the FreeBSD 4.3 KAME/IPSEC MIP is loosely-coupled. KAME IPSEC handles most of the IPSEC-based tunneling. We assume that the KAME IPSEC has been setup, and then run Mobile-IP which merely overrides any IPSEC policy in order to get Mobile-IP functions themselves accomplished.

So the bottom-line question then remains: which is better? It has been said in the past that any "packet filter" or access list mechanism vis-a-vis firewalls may be dangerous, because if the rule set is complex, it is easy to make mistakes. The FreeBSD 4.3 KAME/IPSEC route-based mechanism however is perhaps more esoteric than any possible ACL mechanism. On the other hand, the bottom-line issue here may simply be portability. Our Mobile-IP implementation when ported to Linux will make a very few, reasonable assumptions about IP mobility features needed by an operating system. By comparison the PSU FreeBSD Mobile-IP/IPSEC implementation did not lend itself to simplicity or portability as it made complex assumptions about the host OS IPSEC and routing socket architectures. Our design is more elegant and much simpler.

Summary

In a narrow sense, we do not know of any related work, other than the comparison of the FreeBSD 4.3 version to the Linux OS version as previously presented. In a wider sense, we could consider competing link-layer, and network-layer systems that are somehow targeted at mobile security. Such systems could include 802.11 WEP (link-layer), or other VPN systems like PPTP, which has been fairly well discredited by Bruce Schneier [14]. The critical question is this: Why is IPSEC not chosen by default as the main vehicle for the delivery of end system to border gateway virtual private networks?

IPSEC has major virtues including:

1. It is not specific to any link-layer, and could be used for cellular telephony wireless, or over Ethernet for that matter.
2. It is not link-specific in terms of hop counts. It can easily be multi-hop across the Internet to a remote home security enclave.
3. It has been widely and opened reviewed in the IETF.
4. Over time, it will improve or at least keep up as it was designed for both replacement of its basic cryptographic algorithms, and key exchange algorithms. Thus it is more extensible than fragile algorithms like WEP.

It may be argued that from the layer 1 and layer 2 "IEEE points of view", the IEEE cannot assume that IETF protocols are in use. What would be wrong then with doing nothing? KISS has its virtues and trying to put complex functions like security into firmware or hardware may be best left to layers 3 and above. One might argue that combined IPSEC/Mobile-IP is not a good combination, because perhaps Mobile-IP is not a good idea. There are those who argue that Mobile-IP may perhaps be inefficient or have other problems. It is not our goal here to argue for or against Mobile-IP. One could just as well combine IPSEC with DHCP. DHCP could be authenticated itself, or perhaps protected by

IPSEC in local security domains, and then IPSEC could take care of the two-way tunnels to and from a home security agent. Obviously with DHCP, and unlike with Mobile-IP, IPSEC cannot take advantage of a fixed IP address as a index mechanism because DHCP IP addresses may vary over time or over link reattachments. However, IPSEC provides for this possibility with its dynamic key management protocol called IKE. According to the IPSEC Domain of Interpretation [13], one can simply setup two-way tunnels with IPSEC using dynamic keying and a fixed higher level name a la "user@dnsname", or according to the DOI document, ID USER FQDN. Again, there is no point in avoiding IPSEC.

References

- [1] J. Binkley, and J. Richardson, Security Considerations for Mobility and Firewalls, IETF draft, 1998,
<http://www.cs.pdx.edu/jrb/jrb.papers/firewall/draft.txt>.
- [2] J. Binkley, and W. Trost, Authenticated Ad Hoc Routing at the Link Layer for Mobile Systems, Wireless Networks, Vol. 7, No. 2, pp. 139-145, 2001.
- [3] J. Broch, D.A. Maltz, D.B. Johnson, Y.C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols", Mobicom, Dallas, October 1998, pp. 85-97
- [4] N. Borisov, I. Goldberg, D. Wagner, "Intercepting Mobile Communications: The Insecurity of 802.11", In Proceedings of MobiCom 2001, July 2001.
- [5] P. Ferguson, and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC 2267, IETF, January 1998.
- [6] S. Fluhrer, I. Martin, A. Shamir, "Weaknesses in the key scheduling algorithm of RC4". Eighth Annual Workshop on Selected Areas in Cryptography, August 2001.
- [7] David B. Johnson and David A. Maltz. "Dynamic source routing in ad hoc wireless networks", In Tomasz Imielinski and Henry F. Korth, editors, Mobile Computing, pages 153-181. Kluwer Academic Publishing, 1996.
- [8] KAME IPv6 and IPSEC project,
<http://www.kame.net> , Sept. 21, 2001.
- [9] S. Kent, and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, IETF, November 1998.
- [10] Local and Metropolitan Area Networks, IEEE. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Standard 802.11, 1999.
- [11] J. Moy, "OSPF Version 2", RFC 2328, IETF, 1998.
- [12] C. Perkins, "IP Mobility Support", RFC 2002, IETF, 1996.
- [13] D. Piper, "The Internet IP Security Domain of Interpretation for

ISAKMP", RFC 2407, IETF, 1998.

[14] B. Schneier, and Mudge. "Cryptanalysis of Microsoft's Point-to-Point Tunneling Protocol (PPTP)", 5th ACM Conference on Computer and Communications Security, pp. 132-140, San Francisco, California, November 1998. ACM Press.

[15] Secure Mobile Networks project,
<http://www.cs.pdx.edu/research/SMN> , Sept. 21, 2001.

[16] K. Sklower, "A Tree-Based Packet Routing Table for Berkeley UNIX", Proceedings of the 1991 Winter USENIX Technical Conference, January 1991.

[17] A. Stubblefield, J. Ioannidis, A. Rubin, "Using the Fluhrer, Mantin, and Shamir Attack to Break WEP", ATT Labs Technical Report, TD-4ZCPZZ, Revision 2, August 21, 2001.

[18] T. Ylonen, "SSH - Secure Login Connections over the Internet", USENIX Security Conference VI, 1996, pp. 37-42.

[19] J. Zao, S. Kent, J. Gahm, G. Troxel, M. Condell, P. Helinek, N. Yuan, and I. Castineya. A Public-Key Based Secure Mobile-IP. MobiCom97, September 1997.

APPENDIX F

Secure Architecture for Extensible Mobile Internet Transport Systems

Sixth Quarterly Report

Dr. Digen Das, Dr. Patrick Fitzgibbons, Dr. Larry Hash

15 October 2001

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Accomplishments

In this section, we will briefly present our accomplishments for the past month. We will minimize details and provide only bare bones summary text.

Our accomplishments are discussed in the following subsections.

1. Creation of a secure enclave model for wireless mobility that includes both inter and intra-domain Mobile-IP.

This is a prerequisite for an integrated Mobile-IP/IPSEC system in which Mobile Nodes abroad can use 2-way tunnels to securely tunnel all their packets to and from their Home Agent (assumed to be at home in the secure enclave). It needs to be understood that our solution only addresses one possible facet of a many-sided security problem. We also will need to test various forms of ad hoc routing (including multi-hop) and tie them to end to end (but network-layer) IPSEC-based routing. Thus hosts that a priori belong to the same security enclave may choose to securely communicate to their security peers. Further our Home Agent and Foreign Agents will use one-way tunnels authenticated with AH. This allows all agents to reject any arriving (tunnel) packets that do not have an IPSEC binding using agent IP addresses. We have not overlooked the fact that 2-way tunnels deal with the problem of what to do about mobile nodes, but in so doing they neglect dealing with non-local mobile nodes. Issues here are complicated and we cannot claim to have the final solution. However we have addressed many security issues in this area. In brief, we suggest that foreign nodes simply be logically treated as being "outside" the enclave and their packets need only be tunneled across the enclave to a firewall access point.

2. Preliminary investigation of Mobile-IP and IPSEC in terms of routing and security.

A key focus of our work will be to tie Mobile-IP and IPSEC directly together. At this point in time, Cisco (for example) has made both Mobile-IP and IPSEC available in their routing devices in IOS version 12.0. However there is no evidence that the two have been integrated. We believe that the combination of IPSEC and Mobile-IP is far superior to virtual link-layer tunneling schemes a la L2TP or Microsoft PPTP simply because IPSEC has wide-spread architectural utility, generality, and is liable to be more supported than more proprietary schemes. For example, IPSEC can cover both the link layer (by being at the network layer) and the transport layer OR run router to end system, router-to-router, etc. There is also no point in separate security mechanisms for the latest virtual link-tunneling scheme when IPSEC can be used in all places. Mobile-IP needs IPSEC by definition as all packets between a remote Mobile Node (or Mobile Node on a wireless link) to/from home should be made secure.

In our design, we intend to tie IPSEC to routes. For example, when a Mobile Node installs a default route, it is aiming that route at an agent. A route binding that included an indirection mechanism (tunnel IPSEC to the Home Agent) was part of the picture. When we investigated what was done in the predecessor to Mobile-IP work (e.g., using multi-hop ad hoc routing protocol) to tie two Mobile Nodes together, it became clear that a natural extension was to tie IPSEC to the host routes installed in each Mobile Node. All other IPSEC implementations we have seen so far tie IPSEC to some sort of additional packet-filter like access list mechanism.

Our approach seems more powerful and in some ways simpler, but to be fair we cannot make a compelling case for its superiority over access list mechanisms (other than one less lookup in the IP layer, but even that is not crucial given the relatively blinding speed of

processors these days). Our IPSEC route binding mechanism can however easily be used to setup manual Virtual Private Networks between two routes simply by installing symmetric keys in a key file, and using the route administrative command to install a route (to a network, subnet, or host). Our design also applies to ARP/link-layer bindings. Therefore we believe our architecture has some general utility.

3. Simplified Link-Layer Only Ad Hoc Routing

In this protocol, we do not use ARP on a link. We instead use a protocol (like the ISO ES-IS) in which all nodes, agents, and Mobile Nodes send authenticated beacons. This ARP replacement mechanism is intended to serve a number of purposes. First, it tries to mitigate possible ARP spoofing by insisting that the (IP Address/MAC address) binding is authenticated. Note that our implementation of this mechanism requires both symmetric and asymmetric key systems (in the former case, we have a network-wide key; in the latter, a per host signature). Secondly, the mechanism serves to tie networks together by key possession. It is not important if two laptops do or do not share a subnet. All systems beacon. Therefore if you share a key, you can communicate. The low-level IP subnet semantic that requires a router for communication between two hosts from different subnets is obviated. The mechanism also serves a gateway function so that systems that do not possess the secret cannot penetrate into a secure enclave through a "firewall-like" mobility agent. Lastly the mechanism serves a very important purpose in that we assume that if we can hear your beacon, we can communicate to you. Beacons (unlike ARP) is done at a relatively high rate of speed. If a system disappears, we will use other routing mechanisms to try and find it (and not believe an ARP cache entry that is going to remain static for around twenty minutes). The only drawback of this system is that if everyone beacons the link itself may be less scalable/useable in terms of throughput.

In light of the fact that existing wireless links cannot support many simultaneous hosts anyway, it is hard to understand how this would present a problem. ES-IS originally was criticized along these terms, but the critics apparently did not notice that beacon rates were extremely slow (once per minute for fixed Ethernet systems). Slow beacon rates for Mobile Nodes along with a combined unicast ACK "reply" to agent beacons make the mechanism more scalable. Agents can beacon and Mobile Nodes can simply append their MAC address under Mobile-IP authentication when they use Mobile-IP to register. This takes care of Mobile Nodes that only desire to talk to the wired infrastructure and do not want ad hoc service. Non-mobile Mobile Nodes do not need to beacon very often and can probably slow down their beacon rates (as long as agents do not remove them from the routing state in the agent). Highly Mobile Mobile Nodes need to beacon at higher rates in order to talk to each other.

We suspect the real problem here is wireless loading. If the link itself should be able to tolerate 100 FTP transfers at the same time, one should not worry overly much about 100 nodes sending out 100 byte beacons, even if the beacon rate for all nodes is 1 per second. Of course, this would be too much for WAN wireless systems with small overall bandwidth,

but in such cases an IEEE 802.11 style link-layer registration protocol only between agent and Mobile Node can make sense; i.e., one could well neglect the ad hoc function when there are too many nodes in a cell, or one simply doesn't care to communicate to anything other than the agent (highly likely in many usage scenarios).

4. Multi-hop Multicast Ad Hoc Routing (MADRP)

Our recommended multicast ad hoc routing protocol was an extension of the design based on ideas freely borrowed from early suggestions by Dave Johnson of CMU and Scott Corson of the University of Maryland. This implementation differs from others in that:

1. It does not neglect security between the systems.,
2. It considers redundant communication paths as possibly important
3. It can be integrated with Mobile IP so that it allowed systems more than one hop away from an agent on a wireless link to still communicate with the Internet.

The basic concept is that a source Mobile Node would use a multicast discovery packet to do an expanding ring search for another destination ad hoc host across any number of participating mobile hosts acting as routers. A multicast discovery packet would be sent out and "flooded" or forwarded until it reached either the desired Mobile Node or any Agent. Note that the flooding occurs across systems with only one interface, which is unlike how conventional routing protocols work, as in general, they will take a control packet in one interface and flood it out other interfaces; i.e., Mobile Nodes are assumed to have only one wireless interface. Intermediate systems and the final system would setup host routes pointing back towards the sender. A discovery ACK would be returned to the sender by the receiver which could either be multicast or unicast. The return trip packet would also cause routes to be set up (the send/receive discovery packets always cause routes in the opposite direction to be set up). Thus two-way communication would be enabled. Ordinarily the returned ACK might be unicast back across the "best" path. However if redundancy was desired, one could choose to multicast that packet back, thus informing the source of possible multiple paths to the target.

We also would improve redundancy by considering agents as possible potential default routes; i.e., agents would always claim that they knew the way (even if they didn't). Thus a Mobile Node might get at least two path answers back:

1. From the true Mobile Node target if available, and
2. From a default agent.

The latter would be used only if the former was not available. Our approach would allow Mobile IP to work in the multi-hop case by considering the Home Agent one more remote ad hoc node, and we would search for it. This would produce a host route (through a nearby agent) that would allow Mobile IP to work. We submit this is a rather elegant architectural idea. End to end Mobile

Node security would be implemented by setting IPSEC route bindings up to the source and destination host routes respectively. Thus end path (from source to destination or destination to source) used IPSEC. Note that this mechanism is end to end, not end to router. It does not suffer from the possibility of a proposed plaintext attack. Intermediate systems are presumed to usually not impose additional IPSEC bindings. MADRP routing security itself is similar to Mobile-IP authentication. The protocol name "MADRP" stands for Multicast Ad hoc Demand Routing Protocol. The basic concept is that each participating node would have the same function for metric computation based on a network-wide configurable weight function consisting of $m = f$ (hop count; power remaining; signal strength).

Thus there would be three possible variables in the metric, and weights could be set by the administrator to determine which of the three or combinations therein would be used. For example, one could decide that power at routes with better power and signal-strength would be chosen. Hop count would be ignored or one might setup a link based singly only on power, hop count, or signal-strength. We felt that a combination of signal strength, and power would make the most sense.

5. The Home Agent Redundancy Protocol (HARP)

The SAFEMITS design calls for a protocol that is used between two or more Home Agents to share Mobile registration state (including IPSEC bindings). The agents essentially tunnel in parallel to Mobile Nodes. This mechanism is transparent to Mobile-IP as a whole. Mobile Nodes and Foreign Agents are not a detected. HARP is an Interior Gateway Protocol. We assume that a protocol such OSPF is in use, and that a partitioned subnet scheme will be used. As a result packets sent to HARP agents may be evenly distributed by OSPF due to equal multi-path routing. Our recommendation is that Home Agents not be on the same subnet, so that local power or routing failures cannot take out both Home Agents simultaneously. The downside of HARP is that one must manually administer at least two sets of tables including Mobile-IP registration and IPSEC keys. One might attempt to "simplify" here by introducing yet another single point of failure.

6. The Establishment of Wireless Campus Infrastructure Test bed

We established a limited three node wireless network that has a few users (SAFEMITS project personnel including faculty, and a graduate student) within the School of Information System Engineering Technology at SUNY Institute of Technology at Utica/Rome. We intend to maintain this network and extend it where possible.

Mobile Security Policy Overview

In this section, we will briefly present our overall thoughts on mobile security policy. We will first look at the situation from the point of view of a given secure enclave trying to

implement mobility, and then we will review the situation from the point of view of the traditional Mobile-IP architectural elements (Mobile Node, Home Agent, Foreign Agent)

1. Secure Enclave approach

By "secure enclave" we mean one set of hosts under a single security administration that is protected by some sort of a priori security mechanisms. For example, the secure enclave may be hooked up to the Internet but is protected by one or more firewall systems, which might either, be of the packet filter or bastion host variety. We do not rule out "defense in depth" for interior hosts. We do however assume that some hosts are exposed to the Internet and others are not. We are interested in hosts that are exposed to the Internet. We are also interested in hosts with wireless link layers, which for the sake of argument, we will assume are more susceptible to attacks like promiscuous mode sniffing.

In combining Mobile-IP and IPSEC in Mobile Nodes, we first needed to make number of simplifying assumptions. We assumed that a Mobile Node when at home could maintain a two-way IPSEC tunnel connection between it and its Home Agent. We assumed that a Mobile Node when abroad at either a DHCP link or Foreign Agent link could use two-way IPSEC to tunnel home to the Home Agent. The mobile node would dynamically discover these situations and setup tunnels with appropriate security mechanisms. The Home Agent was thus always a bastion host or security gateway to the secure enclave, via any link on it (wireless or wired). Foreign Agents in our first-cut model only serve as link-layer wireless gateways to a secure enclave and may or may not serve external visitors (but must serve internal wireless users, else they are not of interest). IPSEC associations between Mobile Nodes and Foreign Agents did not exist. A Mobile Node abroad might thus be viewed as an extension of the home secure enclave. The two-way tunnel to and from home would serve as an umbilical code to extend the umbrella of the secure enclave to the Mobile Node.

In more detail, let us consider the relationship between a Mobile Node and Home Agent with the Mobile Node at home. If the Mobile Node is a wireless node, our IPSEC system gives it a functional equivalent for link-layer security. (If it was wired, the user (or the local security authorities) might disable this function). What is important is that all packets barring ARP or Mobile-IP itself would be subject to network-layer IPSEC security, which might be more or less good depending on the specific IPSEC implementation and security transforms in use. Our design calls for replacing ARP with a more secure ad hoc mechanism that simply made traditional ARP spoofing more difficult and made two-way exchanges into the secure enclave (via any agent) difficult as well. Mobile-IP had its own authentication mechanism. There is nothing here to prevent the use of additional security at non-network layers including IPSEC at the transport layer, or transport-equivalent security mechanisms like the Finnish ssh. Here the mechanism is very general. The Mobile Node at home is simply equivalent to any current system using its default router. What is curious is that one still finds weak mechanisms such as the 802.11 WEP (Wired Equivalent

Privacy) in which an algorithm like RC4 may be used for confidentiality between Mobile Node and "bridge" (access point), but is unlikely to be sufficiently strong both because the key length will be restricted due to export reasons, and there is no provision for a key management protocol with the power of protocols in IPSEC. In some respects our design seems to be in competition with 802.11 and yet we believe our overall integrated system is far superior in many ways.

2. "Us versus Them"

When one thinks about a secure enclave, one must consider the enclave from two possible points of view. First one must consider mobile systems that belong to "our side". One must then consider mobile systems that may belong to less trusted visitors. Obviously different security policies might exist for wireless (or wired) mobile systems that belong to the home team as opposed to possible potential visitors (or hackers trying to gain access via a wireless link available via a passing motor vehicle). For example, one might choose to allow local wireless nodes access to the secure enclave and might disallow visitor nodes. Or one might allow visitors access to the wireless network, but disallow access to key internal secure enclave areas. Certainly any number of policies might exist. It seems that exibility in this area could be useful. On the other hand, rich security policy implementations could easily be confusing and hard to get correct.

If we are successful in obtaining additional funding to continue our research for a second year we plan on doing some policy work along these lines, which we anticipate would result in some interesting applications (see below under foreign agents for some discussion of one possible interesting security extension).

This potential follow on research can be summarized in two ways:

1. We would investigate the possibility to limit access via "mobile" link-layer interfaces into the secure enclave. Visitors might be totally disabled or allowed access on a case by case basis. This for example might be done via an ad hoc authentication mechanism. In this instance a mobile node w required to know a shared secret (or present a signed beacon) to an agent. If the agent recognized the beacon, it would install a route to give the mobile node access. If the route was not installed, the mobile node might be able to initiate one-way attacks on the enclave, but it would lack the means for two-way communication. Visitors could be given a temporary key to allow them temporary access into the enclave. We anticipate that while would begin experimenting with ad hoc mobile routing with symmetric keys; we would also experiment with a DNS-based database system for asymmetric keys which would give more key scalability. The critical idea here is that agents act as routers and do not bridge packets naively into the interior infrastructure.
2. We envision that a network design might be employed to deal with the problem of remote visitors (or local wireless systems) by simply rerouting the internal pipes. For

example, any packets coming in on an external unsecured link might simply be tunneled to come back in via an outside external firewall interface. Thus one can easily enable visitors (or local untrusted wireless access) by designing them "outside" the firewall. Packets from trusted external hosts might be allowed direct interior access as long as IPSEC is used (and this risk is deemed worth taking). Mechanisms used here might include known tunnel technologies like Cisco's GRE or IPIP, or even IEEE virtual LANs at the link-layer. The tunnel endpoint (from agent to firewall) must tie to the same sort of input packet filter checks imposed on ordinary Internet packets coming back into the enclave through firewall systems.

3. Foreign Agent Considerations.

Security policy for foreign agents in our original design was simplified and very straightforward. We class foreign agents as either trusted or non-trusted and implemented mechanisms to allow foreign agents to both exclude Mobile Nodes at the link-layer and securely accept tunnel packets from the Home Agent (basically with IPAHIP). Confidentiality was left to the Mobile Node; i.e., the Mobile Node is responsible for making sure that its own packets are secure to/from the Home Agent as it might be using an untrusted Foreign Agent. The reason for using IPSEC authentication in tunnels is to exclude tunnel spoofing possibilities; i.e., the possibility that an attacker might use barebones IPIP to send packets into an infrastructure at an agent, have them decapsulated, and thus appear to be local with a local IP source address. This is not possible if the Foreign Agents only accept IPAHIP packets from Home Agents that they trust and throw away IPIP packets.

In what follows we will briefly present a more complex security policy system (and implementation) below that we may wish to consider if we were to receive continued funding for the SAFEMITS project for a second year. This would allow Foreign Agents and us to consider more complex policies vis-à-vis Mobile Nodes.

4 Home Agent Considerations

Home Agents serve as a bastion-host for mobile systems. Two-way tunnels terminate (or originate) at the Home Agents. It is assumed that barebones (HA to FA) IPIP tunnels are not used with Mobile-IP. Instead one ties IPSEC into the tunnel mechanism. We have already discussed how Foreign Agents could require that all tunneled packets must a priori have some sort of IPSEC association between the Foreign and Home Agent. The Home Agent also serves as the tunnel destination for Mobile Node packets coming back to the enclave. It can enforce a similar semantic; i.e., insist that all inbound IPIP packets must have a Mobile Node/Home Agent (or "our-side" Foreign Agent/Home Agent) security relationship. Barebones unsecured IPIP packets would be disregarded. Thus the Home Agent could defend the security enclave. One downside here is that Home Agents in this system are not end systems; they are intermediate systems (routers). Thus IPSEC packets may be subject to proposed plaintext attacks, as a "man in

the middle" attacker might send packets to the Mobile Node to the Home Agent, and then observe the encrypted packets arriving at the Mobile Node. Defenses against this problem might include session key mechanisms that limit the exposure of keys and/or firewall mechanisms that do not allow Mobile Nodes abroad to communicate to systems that are not in the secure enclave. As a matter of policy, it might be reasonable to assume that Home Agents cannot suffer from a single point of failure scenario. We would likely use the Home Agent Redundancy Protocol (HARP) so that Home Agents could act in parallel. This would need to be accomplished in such a way that IPSEC associations were shared and that in general, Mobile Nodes had no knowledge of HARP.

5. Mobile Node Considerations

In summary, we envision that Mobile Nodes at home would use two-way IPSEC to communicate to the Home Agent when an unsecure link is in use. (Note that this implies that a Home Agent should have an exterior and interior secure enclave interface). When abroad, they should use two-way IPSEC tunnels to both defend against malign influences on less secure links, and/or possible interception across the Internet. We regard concerns about "triangle routing" as irrelevant to security concerns. In general, security between parties who have no trust relationship is an oxymoron. The real security policy considerations that need to be addressed for systems outside the secure enclave are twofold:

1. Should that system be allowed to communicate to home? If the answer is yes, mechanisms such as two-way IPSEC tunnels could be employed.
2. Should systems that are away be allowed to communicate to untrusted systems outside of the secure enclave? If the answer is no, this would obviate such notions as routing redirection targeted to "triangle routing"

We also wish to point out that any integrated solution includes the possibility of so-called "ad hoc" Mobile Nodes engaged in secure communication. With both our ad hoc systems, Mobile Nodes with a priori trust relationship could setup end-to-end IPSEC tunnels and thus securely communicate using legacy protocols like telnet and ftp. These tunnels were at the network layer, but unlike the Mobile Node to Home Agent relationship, they were end to end. Thus it is not possible for any attacker to forward packets through a Mobile Node and create a proposed plaintext attack.

Another possibility we considered is the notion that the "ad hoc" mobile network could be based on shared trust. In this case shared trust might possibly entail something as obvious as ad hoc networks using two shared symmetric keys network-wide. One key would be intended for "our side" and one key intended to be temporarily created and shared with the "others" deemed temporarily trustworthy (and then revoked).

Plans for fourth quarter 2001 and first quarter 2002

In this section, we are going to briefly discuss our short term goals and planning. It is not clear that we will be able provide an actual software prototype in the first year as it has become readily apparent that we must include Linux instead of FreeBSD.

Our plans for an initial release minimally will include the following attributes:

1. An integrated Mobile-IP/IPSEC based on routing where policy is specified very simply in configuration files.
2. Mobile-IP/IPSEC should include the ability to setup IPIP "tunnels" that use either AH/ESP as supplied in the current kernel infrastructure and agents should be able to only accept packets over tunnels that have IPSEC attributes.
4. The release will be targeted towards the Linux operating system.

At this stage we do not choose to predict when we will make this release other than to say it will largely hinge on whether we receive continued funding for an additional year. We find it necessary to point out that we must use what we have and cannot afford to deal with all of the many moving targets at once. The list of possible targets includes IPSEC proper (AH/ESP/combined) including Mobile-IP, and Linux itself. We must first limit the moving targets to switch our platform from FreeBSD to Linux. Of course, we still must integrate other existing subsystems architecturally, which includes more original routing/IPSEC work. For now, we plan on porting our entire system to Linux as soon as possible. In point of fact, that work has already started and we have just begun porting our mobile node to work under Linux.

IPSEC and Mobile-IP - Firewalls and Tunnel Considerations

We are planning to demonstrate that two MNs could simultaneously "borrow" the FA's COA and setup IPIP tunnels back to home systems where the individual tunnels either had or did not have ESP bound to them. We are experimenting with using a script that would determine the COA from existing Mobile-IP utilities and then automatically install the tunnels. There are at least two points worthy of mention:

1. With ESP, all packets are sent from the HA were sent using IPSEC, but we could either bind the default route to send all packets home using IPSEC or leave it alone, and bind IPSEC to the individual tunnel routes that sent packets home. This sort of policy choice is yet another mechanism that we would like to somehow build into our configuration.
2. There is no problem in using the FA's COA since the FA itself (as is probably the case with routers where firewall oriented checking is not on) do not examine

Secure Architecture For Extensible Mobile Internet Transport Systems

First Monthly Report

May 15, 2001

Dr. Digen Das, Dr. Patrick Fitzgibbons and Dr. Larry Hash
SUNY Institute of Technology at Utica/Rome

RE: U.S. Air Force Agreement No. F30602-01-1-0518

Introduction

The Secure Architecture For Extensible Mobile Internet Transport Systems (SAFEMITS) project at SUNY Institute of Technology at Utica/Rome formally began upon receipt of official written notification on April 15, 2001. This being the first of our monthly status reports discusses our activities in setting up an initial test system.

The Initial SAFEMITS Prototype

Current Project Status and Operating Assumptions

The conclusion of the SUNY Institute of Technology Spring semester term was May 5, 2001. As of May 7, 2001 all three of the project investigators began working on the SAFEMITS project on a part time basis. Additionally we have arranged for a graduate student to assist specifically with the software programming aspect. Mr. Amitabh Pandey is an experienced software engineer who has chosen to work on IP network security architecture for his master's project.

We are anticipating that we will be able to undertake at least a two year project. The first year will be as previously planned; i.e., we will implement Mobile-IP and network layer security and integrate them with attention to tunnels. We will use Wavelan as a link-layer medium and will construct a wireless networking infrastructure in a lab located in Room 1196 in Donovan Hall which has been dedicated specifically for this purpose. We have already mapped the room and have determined that a single base station is enough for coverage. We will also attempt to understand and analyze those protocols in a mobile wireless LAN based environment. If we receive funding to continue the project for a second year, the focus will expand to include routing redundancy; i.e., multiple foreign agents, multiple home agents, and "ad hoc" networking; i.e., direct secure communication between mobile hosts.

One question we faced immediately was what operating system to use for a developmental environment. We identified two possibilities as being technically feasible. One option involved using Linux, since Wavelan drivers were already available, or we could use a version of Unix called FreeBSD. At this point we have decided against using

Linux, provided that we can successfully port Wavelan drivers for use with FreeBSD. Our reasoning for opting for FreeBSD was because certain aspects of the current Linux TCP/IP stack are relatively immature. For example, the routing/forwarding code in the Linux kernel consists entirely of a linear search. This precludes finding the longest matching prefix in a subnetted environment. The TCP protocol itself has some questionable areas hence we were more comfortable adapting a more “battle proven” operating system such as FreeBSD. We also felt that the BSD stack is a good choice for other reasons too since there is a widespread basis for sharing with other efforts being done elsewhere. As an example, we are hoping to examine and possibly reuse work being done by Jim Binkley at Portland State University as he has successfully implemented a prototype IPSEC protocols in FreeBSD. Moreover, some work has already been completed to port the Wavelan drivers to FreeBSD. We will require drivers for both ISA bus cards (for wired infrastructure agents; i.e., HA’s and FA’s) and for the PCMCIA cards that are used with the laptops.

We have ordered a first round of equipment consisting of three SONY VAIO Pentium III 750 MHz laptop computers. We have also purchased three VAIO wireless LAN cards and a base station transmitter/receiver.

Our plan for the remainder of the Spring/Summer is to get the link-layer drivers established and then work in the design and implementation of mobile-IP. Our objective for the Fall quarter is to finish Mobile-IP, design and implement network layer security and integrate them. Testing and analysis will most likely continue into the Winter quarter. We also hope to have an implemented protocol deployed in the lab to allow for roaming.

Research Notes

Mobile IP. Mobile-IP has many feature points, perhaps too many as we would like to point out that it is a classic case of “design by committee”. Therefore, we feel we would not be able to make any significant headway by trying to implement them all. As a result, we are attempting to produce a Mobile-IP that, while more limited in terms of features, includes all security features, plus features that we deem essential with respect to security (e.g., two-way tunnels). For example, Mobile-IP allows authentication between all FA/HA/MN pairs and we intend to implement and use that capability. There are two forms of protection permitted and we intend to implement both forms and experiment with them.

What follows are a few examples of the inherent problems with the Mobile-IP semantic set along with our approach for addressing them:

1. Mobile-IP somewhat unclearly assumes that the architecture employs ICMP router advertisements and solicitation (RFC 1256). We do not plan on using solicitations at this point. Agents (FA’s/HA’s) will solicit and we regard the ICMP router advertisements as network-layer beacons, analogous to HELLO messages used in the ISO/OSI ES-IS protocol. Our infrastructure will rely on Foreign Agents as the feasibility of having enough administrative control to

require users to deploy DHCP is not at all a certainty. Further, beacons from agents can provide link-layer signal strength status. This information can be used to improve how and when mobile nodes switch between agents.

2. Mobile-IP permits three different kinds of encapsulation, the basic required IP-IP form, Cisco's GRE, and the so-called minimal encapsulation. We plan on implementing only the basic unicast IP-IP form at the start. The Mobile-IP standard states that the basic form of IP-IP should have a "soft state" mechanism that basically allows a HA to keep track of tunnel state and do something more intelligent with certain ICMP messages returned regarding encapsulation packets. Further, the tunnel should use the PATH MTU mechanism. While we understand the basis for these requirements, it is not our intention to implement "soft state" at this time. We only require bare unicast tunnels between mobile nodes and agents as deemed appropriate. The required tunnel mechanism is complex, given PATH MTU, soft state, etc. There does not appear to be any extensive implementation experience in this regard. Moreover, the standard as written overlooks the fact that unicast tunnels require the exact same transport protocol as used with current multicast IP-IP tunnels (e.g., MBONE). These tunnel mechanisms will need to coexist. We plan on implementing a simple form of IP-IP tunnel that will be accessible as a route to a virtual device in the routing table. Our "bare bones" IP-IP should be, at the very least, interoperable with other implementations.
3. Mobile-IP systems may choose to use one of two forms of protection against "replay" attacks. A timestamp mechanism based on NTP timestamps may be used or a nonce mechanism may be used. For timestamps, the bottom line is that NTP must be used and presumably authenticated possibly from source to mobile node.
4. Mobile-IP optionally permits forwarding of broadcast and multicasts (via a unicast encapsulation scheme) to a remote mobile node. We do not intend to implement these mechanisms.

The above are just a few examples of the issues that we will need to address. It should not be overlooked that Mobile-IP as it stands may contain semantics involving NTP, ICMP router solicitation, PATH MTU, DHCP, multicast routing, and it comes in two basic modes (FA-based and MN sans FA). Additionally, Mobile-IP routing daemons may have to operate alongside of routed, gated, m-routed, and routing discovery. Although the protocol itself is simple, the resulting architectural and test considerations are by no means trivial.

One problem with IETF standards is that in general, a working group might produce a protocol, but there is no guarantee that resulting application/operating system implementations and architectures will be feasible. It is our observation that both Mobile-IP and IPSEC reflect this discontinuity to some extent.

In the case of IPSEC, it is not clear at what layer security associations should be made in the operating system. If we use FreeBSD as our reference, we might decide to tie IPSEC

security associations to sockets (application/transport layer), protocol control blocks (transport layer), or to routes (network layer). The choice may profoundly affect the usability of the system. It is also important for system security and it may turn out that there is in fact no correct choice. For example, both sockets and routes may need security associations.

With Mobile-IP, one problem is that current IP subnet model has to be violated. For example, Foreign Agents must be able to talk to any Mobile Node from any subnet and not pay attention to the MN's subnet home IP address. This is contrary to the way current kernels work as they determine the subnets they are attached to from looking at the IP addresses associated with their interfaces. The BSD stack has another problem which is that it is not possible for an application to determine the input interface for an incoming packet. It is difficult, but not impossible, for applications to direct packets out a given interface. The latter mechanism exists in a crude form and is referred to as a "route to interface". Typically a routing daemon (this only talks to local links anyway) may send a UDP packet to the directed broadcast address (e.g., via subnet .255) and the subnet address can be used to route the packet out the desired interface. Routing daemons that receive packets determine the interface by matching the incoming broadcast subnet with it's own set of interfaces.

Without going into great detail, our approach is to break the design of Mobile-IP down into state machines based on the FA/HA, and/or MN role. There is a state machine model for each. Inputs are time timeouts or ICMP or UDP packets. Outputs are actions such as inserting or deleting routes, dropping a node from a visitor list, and/or sending control packets. The outcome of this decomposition is that we believe we can successfully separate the functionality between application/routing daemons that will implement policy and a select number of bare-bones kernel mechanisms. It is obviously desirable that as much functionality as possible should be designed into the application layer and as little functionality as possible in the kernel.

Ad Hoc routing protocols. We have conducted a literature search on the subject of "ad hoc" routing protocols. We will report on this in more detail in the next status report and also include the appropriate references. At the present time it will suffice to say that there is a limited body of highly theoretical work on the subject and in general, it can be divided along two entities. The first entity is attempting to reuse existing routing protocols (distance-vector and/or link-state) even though the need to minimize broadcast and involvement of every mobile node is by means a trivial achievement. The second entity as embodied by those researchers at Carnegie Mellon University under the direction of Dave Johnson, proposes that something might be learned from previous radio efforts that it is conceivable that a routing protocol based on demand source routing might minimize the involvement of other hosts. A demand source routing system coupled with network authentication for security could provide an intriguing research alternative.

Link-layer routing and security observations. We have identified two possible physical networking configurations that could be employed for constructing a Whelan-

based LAN (not a WAN). The constant here is that in either instance we will use PCMCIA cards in laptops for mobility. What varies is how we arrange for connectivity to a wired infrastructure. We could choose to either: 1) employ ISA cards in PCs that would function as either bridges or routers, or 2) purchase an off the shelf “Access Point” which is essentially an Intel Pentium based computer that comes pre-packaged with an Ethernet card and a Wavelan ISA card installed. The interesting difference is that the Access Point comes with what the vendor terms “roaming” software.

The vendors of such products as Wavelan are limited to what they can do to enable mobility and can only operate at the link or physical layers (the domain of device drivers). They are limited to “same network” in terms of the network layer; e.g., IP routing will fail if the transmitting node moved from one base station to another and the base station was on a different IP network segment. Mobile LAN vendors by definition have to incorporate a bridging solution. It is also desirable that the number of packets on the radio link be minimized. One partial solution is referred to as “roaming” which effectively means that a link-layer registration protocol is used between PCMCIA-based clients and base stations (Access Points). The latter is a software solution. AP’s send “beacons” which are periodic special link-layer packets that have a special header on top of the normal link-layer header. The mobile node (client) uses a link-layer sign-on protocol in response to any received beacon that has the higher signal level. As a result, packet duplication is eliminated. AP’s that receive packets destined for a given mobile node can discard them if the node is not registered with them. This may be a useful optimization, although it does not address network layer mobility. There may also be cause for criticism on a couple of points with respect to secure networking.

Link-layer beaming may have an inherent flaw in that an outside attacker could for example pull up in a vehicle and run an Access Point that might have a better signal strength and thus entice all packets to be transmitted in its direction. This is in some sense a general problem with any routing protocol with a centralized system that requires that “all packets be sent to a particular target”. For example, this is also true for the Mobile-IP Home Agent (HA) since tunnels to and from it might be manipulated by attackers or the HA itself might be spoofed. It may be better to simply beacon at the network layer and thus take advantage of IPSEC mechanisms such as authentication. In general, routing control (or all packets) would be much more secure with network-layer authentication mechanisms. For example, ARP could stand to be authenticated as well, although possibly the mechanism here could better be addressed by adapting the ICMP registration protocol for ad hoc systems.

From a redundancy point of view, it may be that signing-on with one system and only using it for access to the wired infrastructure is not necessarily a better alternative. If the link-layer is under stress (packet loss and corruption), which is certainly a likely attribute of radio transmission by definition, multiple paths may be useful.

Bibliography

- [1] Jim Binkley. Authenticated ad hoc routing at the link layer for mobile systems. Technical Report 96-3, Portland State University, Computer Science, 1996.
- [2] Jim Binkley and John Richardson. Security considerations for mobility and firewalls. Internet-Draft, November 1998.
- [3] Josh Broch, David Johnson, and David Maltz. "the dynamic source routing protocol for mobile ad hoc networks". MANET IETF draft, draft-ietf-manet-dsr-01.txt, December 1998.
- [4] Bjorn Chambless and Jim Binkley. Harp - home agent redundancy protocol". Internet Draft, October 1997.
- [5] <http://www.cisco.com>. Internet web site. Search on IPSEC and Mobile-IP.
- [6] <http://www.cisco.com>. Internet web site. Search on L2TP.
- [7] Counterpane systems announces crack of microsoft's point-to-point tunneling protocol. <http://www.counterpane.com/ppptp.html>. Bruce Schneier's Cryptanalysis paper should be read first.
More information can be found at <http://www.microsoft.com>.
- [8] Wireless lan medium access control (mac) and physical layer (phy) specifications. IEEE Std. 802.11b-1999, November 1999.
- [9] linux ieee wavelan driver. <http://www.fast.fh-dortmund.e/users/andy/wvlan>.
- [10] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, first monthly report. May 2001
- [11] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, second monthly report. June 2001
- [12] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, third monthly report. July 2001
- [13] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, fourth monthly report. August 2001
- [14] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, fifth monthly report. September 2001
- [15] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, sixth monthly report. October 2001
- [16] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, seventh monthly report. November 2001
- [17] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, eighth monthly report. December 2001
- [18] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture For Extensible Mobile Internet System, ninth monthly report. January 2002
- [19] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture for Extensible Mobile Internet System, tenth monthly report, February 2002
- [20] Digen Das, Patrick Fitzgibbons and Larry Hash. Secure Architecture for Extensible Mobile Internet System, eleventh monthly report, March 2002